

# Information-Theoretic Lower Bounds on the Storage Cost of Shared Memory Emulation

Viveck R. Cadambe  
EE Department,  
Pennsylvania State University,  
University Park, PA, USA  
viveck@engr.psu.edu

Zhiying Wang  
CPCC Center  
University of California, Irvine  
Irvine, CA, USA  
zhiying@uci.edu

Nancy Lynch  
Department of EECS  
MIT  
Cambridge, MA, USA  
lynch@csail.mit.edu

## Abstract

The focus of this paper is to understand storage costs of emulating an atomic shared memory over an asynchronous, distributed message passing system. Previous literature has developed several shared memory emulation algorithms based on replication and erasure coding techniques, and analyzed the storage costs of the proposed algorithms. In this paper, we present information-theoretic lower bounds on the storage costs incurred by shared memory emulation algorithms. Our storage cost lower bounds are universally applicable, that is, we make no assumption on the structure of the algorithm or the method of encoding the data.

We consider an arbitrary algorithm  $A$  that implements an atomic multi-writer single-reader (MWSR) shared memory variable whose values come from a finite set  $\mathcal{V}$  over a system of  $N$  servers connected by point-to-point asynchronous links. We require that in every fair execution of algorithm  $A$  where the number of server failures is smaller than a parameter  $f$ , every operation invoked at a non-failing client terminates. We define the storage cost of a server in algorithm  $A$  as the logarithm (to base 2) of number of states it can take on; the total-storage cost of algorithm  $A$  is the sum of the storage cost of all servers. We develop three lower bounds on the storage cost of algorithm  $A$ .

- In our first lower bound, we show that if algorithm  $A$  does not use server gossip, then the total storage cost is lower bounded by  $2^{\frac{N}{N-f+1}} \log_2 |\mathcal{V}| - o(\log_2 |\mathcal{V}|)$ .
- In our second lower bound we show that the total storage cost is at least  $2^{\frac{N}{N-f+2}} \log_2 |\mathcal{V}| - o(\log_2 |\mathcal{V}|)$  even if the algorithm uses server gossip.
- In our third lower bound, we consider algorithms where the write protocol sends information about the value in at most one phase. For such algorithms, we show that the total storage cost is at least  $\nu^* \frac{N}{N-f+\nu^*-1} \log_2 (|\mathcal{V}|) - o(\log_2 (|\mathcal{V}|))$ , where  $\nu^*$  is the minimum of  $f+1$  and the number of active write operations of an execution.

Our first and second lower bounds are approximately twice as strong as the previously known bound of  $\frac{N}{N-f} \log_2 |\mathcal{V}|$ . Furthermore, our first two lower bounds apply even for regular, single-writer single-reader (SWSR) shared memory emulation algorithms. Our third lower bound is much larger than our first and second lower bounds, although it is applicable to a smaller class of algorithms where the write protocol has certain restrictions. In particular, our third bound is comparable to the storage cost achieved by most shared memory emulation algorithms in the literature, which naturally fall under the class of algorithms studied. Our proof ideas are inspired by recent results in coding theory.

# 1 Introduction

The emulation of a consistent, fault-tolerant read-write shared memory in a distributed, asynchronous, storage system has been an active area of research in distributed computing theory. In their celebrated paper [3], Attiya, Bar-Noy, and Dolev devised a fault-tolerant algorithm for emulating a shared memory that achieves atomic consistency (linearizability) [16,17]. Consider a distributed system with server nodes, write client and read client nodes, all of which are connected by point-to-point asynchronous links. The ideas of [3] can be used to design server, write and read protocols that implement an atomic shared memory even if the write and read operations are invoked concurrently with the following guarantee: every read or write operation invoked at a non-failing client terminates so long as the set of servers that fail is restricted to a minority. The algorithm of [3] used a replication-based storage scheme at the servers to attain fault tolerance. Following [3], several papers [1, 2, 4, 5, 11, 12, 15, 21] have developed algorithms that use *erasure coding* instead of replication for fault tolerance, with the goal of improving upon the storage efficiency of [3].

In erasure coding<sup>1</sup> which is studied in classical coding theory, each server stores a function of the value called a codeword symbol. A decoder that is able to access a sufficient number of codeword symbols recovers the value. The number of bits used to represent a codeword symbol is typically much smaller than the number of bits used to represent the value. As a consequence, erasure coding is well known to lead to smaller storage costs as compared to replication in the classical coding-theoretic set-up (See, for example, [7,10,19]). Here, we aim to understand storage costs of shared memory emulation, where in contrast with the classical coding-theoretic setup, multiple versions of the data object are to be stored in a consistent manner.

When erasure coding is used for shared memory emulation, new challenges arise. Since, in erasure coding, each server stores a codeword symbol and not the entire value, a read operation has to obtain a sufficient number of codeword symbols to decode the value being stored. When a write operation begins to write a new version of the data object, the old version cannot be deleted from the servers until a sufficient number of codeword symbols corresponding to the new version have been propagated to the servers. As a consequence, servers have to store codeword symbols corresponding to multiple versions of the data object to ensure that a reader can decode an atomically consistent version. Previous erasure coding based shared memory emulation algorithms [1,4,5,11,12,15,21] have noted that the number of versions to be stored at a server can be large if there are a large number of ongoing or failed write operations whose codeword symbols have not been propagated sufficiently. Because servers store codeword symbols corresponding to multiple versions, the storage cost of using erasure coding can be large, even if the number of bits in each codeword symbol is small compared to the number of bits used to represent the value.

Despite the vast amount of literature in the study of storage costs of shared memory emulation, some compelling and fundamental questions remain unanswered. Since a server can store an arbitrary function of all the symbols it receives, can we develop a sophisticated storage strategy that somehow compresses multiple versions at the servers and thereby results in smaller storage costs? If we add multiple phases to read and write protocols or include other algorithmic novelties, can we reduce the storage cost of shared memory emulation? In our paper, we obtain insights into these questions by developing novel impossibility results that lower bound the storage cost of an arbitrary atomic shared memory emulation algorithm.

## 2 Summary of Results and Comparisons with Related Work

In this section, we first summarize the shared memory emulation and the classical coding theory set-ups. We then describe our storage cost lower bounds in Theorems 4.1 and 5.1. Then, we describe our storage cost lower bound related to Theorem 6.5. Finally we compare our results to previously derived storage cost lower bounds.

### 2.1 Set up

**Shared Memory Emulation Set-up:** We consider an arbitrary algorithm  $A$  that implements, over a network of  $N$  servers connected by point-to-point asynchronous links, an atomic multi-writer single-reader (MWSR) shared memory variable whose values come from a finite set  $\mathcal{V}$ . The algorithm  $A$  is required to ensure that all operations terminate so long as the number of server failures is no larger than a parameter

---

<sup>1</sup>Server failures are modeled as erasures of codeword symbols; hence the term, erasure coding.

$f$ . The storage cost of a server in algorithm  $A$  is measured as the logarithm of the number of possible states of the server, and the storage cost of algorithm  $A$  is the total storage cost over all the servers.

**Classical Coding Theory Set-up:** Consider a system with  $N$  servers, where a single version of a data object whose values come from a finite set  $\mathcal{V}$  is to be stored. The value of the data object must be recoverable, so long as the number of server failures is no larger than a parameter  $f$ . The classical *Singleton bound* [18,20] in coding theory implies that the total storage cost is at least  $\frac{N \log_2 |\mathcal{V}|}{N-f}$  bits<sup>2</sup>. The lower bound of  $\frac{N \log_2 |\mathcal{V}|}{N-f}$  on the storage cost is known to be tight in the classical coding-theoretic set-up for large values of  $|\mathcal{V}|$  [18,20].

The power of erasure coding is transparent when we want to design a storage system that tolerates failures of  $f$  server nodes and the number of server nodes  $N$  can be chosen freely. If we use replication, every server stores  $\log_2 |\mathcal{V}|$  bits. Since we need at least  $f+1$  servers to tolerate  $f$  server failures, the total storage cost of the system is at least  $(f+1) \log_2 |\mathcal{V}|$  bits. In contrast, if we use erasure coding, the total storage cost of the system is  $\frac{N \log_2 |\mathcal{V}|}{N-f}$ , which approaches  $\log_2 |\mathcal{V}|$  as  $N$  increases. If  $N$  is sufficiently large, the storage cost of replication is approximately  $f+1$  times the storage cost of erasure coding.

## 2.2 Motivation and Summary - Theorems 4.1 and 5.1

**Motivation:** The classical coding-theoretic model does not model clients or channels, and therefore differs significantly from the shared memory emulation model. However, the storage cost lower bound of  $\frac{\log_2 |\mathcal{V}|}{N-f}$  described by the Singleton bound is, in fact, applicable in the context of shared memory emulation as well. We provide the first formal proof of the lower bound in Appendix B; in particular, we show that for any SWSR regular shared memory emulation algorithm that implements a read write data object whose values come from a set  $\mathcal{V}$ , the total storage cost is at least  $\frac{N \log_2 |\mathcal{V}|}{N-f}$ . The natural lower bound of  $\frac{N \log_2 |\mathcal{V}|}{N-f}$  inspires the following question.

**Question 1:** Does there exist an atomic shared memory emulation algorithm whose storage cost is equal to  $\frac{N}{N-f} \log_2 |\mathcal{V}|$ ?

**Summary of Theorems 4.1 and 5.1:** In this paper, we answer the above question in the negative by proving storage cost lower bounds that are stronger than  $\frac{N}{N-f} \log_2 |\mathcal{V}|$ . In Theorems 4.1 and 5.1 we show that the total-storage cost of single-writer-single-reader (SWSR) *regular* shared memory emulation algorithm is at least  $\frac{2N}{N-f+2} \log_2 |\mathcal{V}| - o(\log_2 |\mathcal{V}|)$ . In particular, if  $f$  is fixed and  $N$  is chosen freely, the total-storage cost lower bound of Theorems 4.1 and 5.1 approach  $\frac{2N}{N-f} \log_2 |\mathcal{V}| - o(|\log_2 |\mathcal{V}||)$  as  $N$  increases; therefore the bounds of Theorems 4.1 and 5.1 are twice as large as the previously known lower bound. Recall that regularity [17] is a weaker consistency model as compared with atomicity. Since Theorems 4.1 and 5.1 apply for regular SWSR shared memory emulation algorithm, they automatically apply for atomic MWSR shared memory emulation algorithms. Theorem 4.1 describe a storage cost lower bound for algorithms which do not use server gossip. Theorem 5.1 describes a lower bound for *any* shared memory emulation algorithm, including algorithms that use server gossip. Our storage cost lower bounds are universal in nature, that is, we make no assumption on the structure of the protocols or the method of data storage. Because we answer Question 1 in the negative, an important implication is that there is an unavoidable price, in terms of storage cost, to ensure regularity in a shared memory emulation system. We next discuss the tightness of Theorems 4.1 and 5.1 in the context of previously derived storage cost upper bounds.

## 2.3 Motivation and Summary - Theorem 6.5

In the sequel, we define the number of active write operations at point  $P$  of an execution as the number of write operations which have begun before the point  $P$  but not yet terminated or failed at point  $P$ . The number of active write operations of an execution is the supremum, over all points of the execution, of the number of active write operations at the points of the execution.

**Motivation:** There is a growing body of literature related to erasure coding based shared memory emulation algorithms [1,2,4–6,11,12,15,21,23]. These algorithms differ in their structure, liveness conditions on operation termination, and their communication costs. A common insight that applies to all the algorithms

---

<sup>2</sup>For the sake of the discussion here, we assume that  $|\mathcal{V}|$  is a power of 2. We refer the reader to [6,18] for more details about erasure coding.

of [1, 2, 4–6, 11, 12, 15, 21, 23] is that, among the class of all executions with at most  $\nu$  active write operations, the worst case storage cost of implementing an atomic shared memory object whose values come from a finite set  $\mathcal{V}$  is *at least*  $\nu \frac{N \log_2 |\mathcal{V}|}{N-f}$ . In fact, references [2, 4, 5, 12] conduct a formal analysis of the incurred storage cost and show that the storage cost incurred<sup>3</sup> is approximately  $\nu \frac{N \log_2 |\mathcal{V}|}{N-f}$ . While the prior works highlight the benefit of erasure coding when the number of active writes is small, the storage cost benefits of erasure coding vanish as the number of active writes increases. In particular, for a sufficiently large value of  $\nu$ , erasure coding based algorithms can even have a higher storage cost as compared to replication based algorithms [3, 13], which incur a storage cost of  $\Theta(f) \log_2 |\mathcal{V}|$  irrespective of the number of active writes.

In contrast with the storage cost upper bounds in literature, our lower bounds of Theorem 4.1 and 5.1 do not depend on the number of active writes. Furthermore, if  $f$  is proportional to  $N$ , then storage cost lower bounds of Theorem 4.1 and 5.1 are both  $o(f) \log_2 |\mathcal{V}| + o(\log_2 |\mathcal{V}|)$ . Prior literature in conjunction with our results of Theorems 4.1 and 5.1 motivates the following question:

**Question 2:** Can we develop an algorithm whose storage cost, when  $f$  is proportional to  $N$ , is as small as  $o(f) \log_2 |\mathcal{V}|$  and does not grow with the number of active writes?

**Summary of Theorem 6.5:** We provide partial answer to Question 2 in our lower bound presented in Theorem 6.5. The lower bound states that the answer to Question 2 is negative, if the write protocol of the algorithm satisfies certain technical conditions described in Section 6. Informally speaking, the technical conditions in Section 6 imply that the write operation is executed in phases, and a message containing information about the value is sent to the servers in at most one phase per write operation. For any atomic MWSR algorithm that ensures that all operations terminate in every execution where the active number of write operations is at most  $\nu$  and the number of server failures is at most  $f$ , Theorem 6.5 shows that if the write protocol satisfies the conditions stated in Section 6, then the storage cost cannot be smaller than  $\nu^* \frac{\log_2 |\mathcal{V}|}{N-f+\nu^*-1} - o(\log |\mathcal{V}|)$ , where  $\nu^*$  is the minimum of  $\{f+1, \nu\}$ .

Theorem 6.5 is interesting from a conceptual viewpoint since it captures the dependence of the storage cost on the degree of concurrency that has been noticed in the upper bounds of [2, 4–6, 12]. In particular, the bound of Theorem 6.5 can be much larger than the bounds of Theorems 4.1 and 5.1, if the parameters  $\nu$  and  $f$  are sufficiently large. If the number of active write operations exceeds  $(f+1)$ , then our storage cost lower bound of Theorem 6.5, which equals  $(f+1) \log_2 |\mathcal{V}| - o(\log_2 |\mathcal{V}|)$ , implies that replication based algorithms are approximately optimal in the class of algorithms described in the theorem.

The class of algorithms that satisfy the conditions stated in Section 6 include a majority of the algorithms in literature [1, 4, 5, 11, 12, 21]. We refer the reader to Section 6 for a more detailed justification. Theorem 6.5, in the stated form, does not apply to a few algorithms [2, 15] because these protocols send messages related to the value of the write operation in two phases; one phase is used to send a hash of the value for client verification purposes, and a second phase is used to send codeword symbols corresponding to the value. In related discussions in Section 6, we conjecture that the statement of Theorem 6.5 and the proof can be modified, without deviating too much from our approach, to apply to a larger class of algorithms which include [2, 15].

In Figure 1 we compare the proposed total-storage lower bounds with the previous achievable upper bounds.

## 2.4 Comparison with Prior Storage Cost Lower Bounds

We compare our work with results of [8, 9, 13, 23–25] which provide some impossibility results in connection to consistent shared memory emulation. The main reference that directly pertains to our work here is [23], which describes interesting, non-trivial lower bounds on shared memory emulation algorithms where the server and client storage schemes satisfy certain restrictions. Reference [23] assumed that every bit stored in the system is associated uniquely with a write operation, and showed that under such a storage scheme, the worst case total storage cost of the system is at least  $\Omega(\min(f, \nu) \log |\mathcal{V}|)$ . The implication of [23] is that in the worst case, if the degree of concurrency is infinite and the server storage scheme is restricted in a particular manner, then the replication based algorithms of [3, 13] are approximately optimal.

---

<sup>3</sup>There are subtle differences in the storage cost incurred by the algorithms of [2, 4–6, 12]. Nonetheless,  $\nu \frac{N}{N-f} \log_2 |\mathcal{V}|$  is a lower bound on the cost incurred by these algorithms in the *worst case*, among executions where the number of active writes is bounded by  $\nu$ .




Figure 1: Storage cost upper and lower bounds for  $N = 21$  servers and  $f = 10$  server failures. We plot the total-storage cost normalized by  $\log_2 |\mathcal{V}|$  when  $|\mathcal{V}| \rightarrow \infty$ . Theorems 5.1, 6.5, B.1 are lower bounds obtained in this paper, that corresponds to  $2\frac{N}{N-f+2}$ ,  $\nu^* \frac{N}{N-f+\nu^*-1}$ ,  $\frac{N}{N-f}$ , respectively,  $\nu^* = \min(f+1, \nu)$ . And ABD and erasure-coding refer to upper bounds achieved in [3] and [2, 4, 5, 12], respectively, which corresponds to  $f+1$  and  $\nu \frac{N}{N-f}$ .

The assumption of [23] that every bit stored is associated with a unique write operation is restrictive and does not apply to all possible storage methods. To see this, consider a scenario where  $\mathcal{V}$  is a finite field. Let  $v_1, v_2 \in \mathcal{V}$  be values corresponding to two different write operations. Suppose in some algorithm  $A$ , at some point of an execution, a server stores  $v_1 + v_2$ , where  $+$  denotes the addition operator over the field. Then a bit stored by the server cannot be uniquely associated with any of the write operations in an unambiguous way. Therefore, the proof technique and the result of [23] fails to provide any insight on the storage cost of the algorithm  $A$ . Put differently, the storage method of [23] does not allow for server storage techniques that potentially compress the values of different versions together (See Appendix A for more technical details). In contrast, the results of Theorems 4.1, 5.1 are universal and would automatically apply to algorithm  $A$ . The result of Theorem 6.5 does not impose any structure on the storage method, and could also apply to algorithm  $A$  if its write protocol satisfies the appropriate restrictions.

References [8, 13] describe impossibility results which are peripherally related to our work. In particular, the results of [8, 13] show that if the readers or writers do not help write another client's value [13], or if the servers are modeled as read-write objects [8], the number of servers must be at least linear in the degree of concurrency in the system. However, the results of [8, 13] do not directly relate to the total-storage cost incurred by the algorithm. Reference [9] considered algorithms where the readers do not change the state of the servers, that is, the readers do not write any values or metadata. The reference showed that for the class of algorithms considered, if the value comes from an *infinite* set, then there exists no regular shared memory emulation algorithm that tolerates even a single server failure. The reference nonetheless does not provide any insight into the storage cost, particularly when the values come from a finite domain.

In information theory literature, recently developed formulations generalize the classical erasure coding model with the goal of understanding storage costs in systems where consistency is important [24, 25]. The models of [24, 25] however, differs from the model considered here. In particular, the models of [24, 25] does not involve formal notions of write and client protocols, and the decoding requirement is only loosely based on the notion of atomicity. Our proofs of Theorem 4.1, 5.1 and 6.5 bear resemblance to storage cost lower bounds of [24].

In our concluding section, Section 7, we provide a summary of the state of the art based on the main results of our paper and of [23].

### 3 Preliminaries

We study the emulation of a shared atomic memory in an asynchronous message passing network. Our setting consists of a set of  $N$  server nodes and a possibly infinite set of client nodes. Without loss of generality, we let the set of server nodes be  $\{1, 2, \dots, N\}$ . We denote the set of client nodes as  $\mathcal{C}$ . We assume a multi-writer single-reader<sup>4</sup> setting, that is, we assume that  $\mathcal{C}$  has a single write client; the remaining clients in  $\mathcal{C}$  are read clients. Each client node is connected to all the server nodes, and the servers are connected to each other via point-to-point reliable, asynchronous, channels. We assume that the readers receive read requests (invocations) from some external source and respond with object values. We assume that writers receive write-requests and respond with acknowledgements. Every new invocation at a client waits for a response of a preceding invocation at the same client. The goal of a shared memory emulation algorithm  $A$  studied in this paper is to design the client and server protocols that implement a read-write register of a data object which can take values from a finite set  $\mathcal{V}$ , with the following safety and liveness properties.

*Safety Properties:* The algorithm must emulate a SWSR regular registers [17]. Informally a regular shared memory object requires that every read operation returns either the value written by the latest write operation that terminates before the invocation of the read operation, or the value of a write operation that overlaps with the read operation. In Section 6 we consider multi-writer single-reader algorithms, and we require the algorithm to be atomic [17]<sup>5</sup>. Informally, in an atomic algorithm, the observed external behavior of every execution looks like the execution of a serial variable type. Recall that SWSR execution of an atomic shared memory emulation is also regular, so our lower bounds for regular algorithms in Theorems 4.1 and 5.1 also apply to atomic emulation algorithms.

*Liveness Properties:* An operation of a non-failed client must terminate in a fair execution, so long as some conditions are satisfied in the execution. Specifically, we require operations to terminate if the number of server failures in the execution is bounded by a parameter  $f$ . We consider algorithms with weaker liveness properties as well in Section 6, where we ensure termination of operations in executions if the number of *active* write operations is bounded. A formal statement of the weaker liveness properties is provided in Section 6.

We require the above correctness properties to hold irrespective of the number of client failures. The data object can take values from a finite set  $\mathcal{V}$ .

### Storage Cost Definition

Informally speaking, the *storage cost* of an algorithm is the total number of bits stored by the servers. In general, for an algorithm where the state of server node  $i \in \{1, 2, \dots, N\}$  can take values from a set  $\mathcal{S}_i$ , we define the storage cost of the server to be equal to  $\log_2 |\mathcal{S}_i|$  bits. The *max-storage cost* of the algorithm  $A$  is defined to be

$$\text{MaxStorage}(A) = \max_{i \in \{1, 2, \dots, N\}} \log_2 |\mathcal{S}_i|.$$

The *total-storage cost* of the algorithm is defined to be

$$\text{TotalStorage}(A) = \sum_{i=1}^N \log_2 |\mathcal{S}_i|.$$

## 4 Storage Cost Lower Bound for Algorithms Without Gossip

In Appendix B we provide a simple but non-trivial proof of the storage cost lower bound that is analogous to Singleton bound. Some of the proof techniques there are also applied in this section. The readers can

---

<sup>4</sup>The storage cost lower bounds presented in our paper apply trivially to multi-writer single-reader shared memory emulation algorithms as well.

<sup>5</sup>In fact, we will study weakly regular multi-writer single-reader algorithms [22]. See details in Section 6

first read Appendix B as an warm-up exercise.

Our main result of this section is a storage cost lower bound, assuming that servers do not gossip. Specifically, in this section, we assume that every message is sent from a client to a server, or from a server to a client. The lower bound is an implication of Theorem 4.1, which describes constraints on the cardinalities of the server states that must be satisfied by any atomic shared memory emulation algorithm where servers do no gossip. The lower bounds on the max- and total-storage costs are stated in Corollary 4.2. After stating Theorem 4.1 and Corollary 4.2, we provide an informal description of the proof of Theorem 4.1, followed by a formal description.

#### 4.1 Statement of Theorem 4.1

**Theorem 4.1.** *Let  $A$  be a single-writer-single-reader shared memory emulation algorithm that implements a regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose that in  $A$ , every message is sent from a server to a client, or from a client to a server. Also, suppose that every server's state belongs to a set  $\mathcal{S}$  in algorithm  $A$ .*

*Suppose that the algorithm  $A$  satisfies the following liveness property: In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$ ,  $f \geq 2$ , then every operation invoked at a non-failing client terminates.*

*Then, for every subset  $\mathcal{N} \subset \{1, 2, \dots, N\}$  where  $|\mathcal{N}| = N - f$ ,*

$$\sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_n| + \max_{n \in \mathcal{N}} \log_2 |\mathcal{S}_n| \geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f).$$

**Corollary 4.2.** *Let  $A$  be a single-writer-single-reader shared memory emulation algorithm that implements a regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose that in  $A$ , every message is sent from a server to a client, or from a client to a server. Also, suppose that every server's state belongs to a set  $\mathcal{S}$  in algorithm  $A$ .*

*Suppose that the algorithm  $A$  satisfies the following liveness property: In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$ ,  $f \geq 2$ , then every operation invoked at a non-failing client terminates. Then*

$$\text{MaxStorage}(A) \geq \frac{\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f)}{N - f + 1},$$

*and*

$$\text{TotalStorage}(A) \geq \frac{N(\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f))}{N - f + 1}.$$

*Proof of Corollary 4.2.* We assume, without loss of generality, that  $|\mathcal{S}_1| \leq |\mathcal{S}_2| \leq \dots \leq |\mathcal{S}_N|$ . From Theorem 4.1, we have

$$\sum_{n=1}^{N-f} \log_2 |\mathcal{S}_n| + \log_2 |\mathcal{S}_{N-f}| \geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f).$$

As a consequence, we have  $\log_2 |\mathcal{S}_{N-f}| \geq \frac{\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f)}{N - f + 1}$ . Therefore, we have  $\max_{n \in \{1, 2, \dots, N\}} \log_2 |\mathcal{S}_n| \geq \frac{\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f)}{N - f + 1}$ . Furthermore, we have  $\log_2 |\mathcal{S}_n| \geq \frac{\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f)}{N - f + 1}$  for ev-

ery  $n \in \{N - f + 1, \dots, N\}$ . This implies the following chain of relations.

$$\begin{aligned}
\sum_{n=1}^N \log_2 |\mathcal{S}_n| &\geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f) - \log_2 |\mathcal{S}_{N-f}| + \sum_{n=N-f+1}^N \log_2 |\mathcal{S}_n| \\
&\geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f) + \sum_{n=N-f+2}^N \log_2 |\mathcal{S}_n| \\
&\geq (\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f)) \left(1 + \frac{f-1}{N-f+1}\right) \\
&= N \left( \frac{\log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f)}{N - f + 1} \right)
\end{aligned}$$

This completes the proof.  $\square$

## 4.2 Informal Proof Sketch of Theorem 4.1

Informally speaking, our lower bound argument is as follows. For every subset  $\mathcal{N} \subset \{1, 2, \dots, N\}$  where  $|\mathcal{N}| = N - f$ , we construct an execution where the servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  fail at the beginning of the execution. The execution has two write operations for values  $v_1$  and  $v_2$ , where  $v_1 \neq v_2$ . The second write operation which writes value  $v_2$  begins after the termination of the first write operation, which has value  $v_1$ .

In this execution, after the point of termination of the first write, a reader can return  $v_1$  because of regularity. Similarly, after the termination of the second write operation, a reader can return  $v_2$ . Therefore, the value  $v_1$  is returnable from the servers at a point before the invocation of the second write operation and  $v_2$  is returnable from the servers at a point after the completion of second write operation. Furthermore, at every point in the interval of the second write operation, at least one of  $v_1$  or  $v_2$  are returnable. This implies that, in the interval of the second write operation, there are two consecutive points  $P$  and  $P'$  such that  $v_1$  must be returnable from the non-failing servers at point  $P$  and  $v_2$  must be returnable from the non-failing servers at point  $P'$ . Since  $(v_1, v_2)$  can be any ordered pair of distinct values from  $\mathcal{V}$ , there must be a one-to-one mapping between the set  $\{(v_1, v_2) : (v_1, v_2) \in \mathcal{V} \times \mathcal{V}, v_1 \neq v_2\}$  and the set of possible configurations of server states at points  $P$  and  $P'$ . This implies that the number of possible server states at points  $P$  and  $P'$  is at least  $(|\mathcal{V}|)(|\mathcal{V}| - 1)$ . Since  $P$  and  $P'$  are consecutive, at most one non-failing server changes its state between these two points. At least  $N - f - 1$  non-failing servers have the same state at point  $P$  as at point  $P'$ . We use this fact to show that the number of elements in the set of possible server states at two consecutive points is at most  $\prod_{n \in \mathcal{N}} |\mathcal{S}_i| \times \max_{n \in \mathcal{N}} |\mathcal{S}_n| \times (N - f)$ . Therefore, we get  $\prod_{n \in \mathcal{N}} |\mathcal{S}_i| \times \max_{n \in \mathcal{N}} |\mathcal{S}_n| \times (N - f) \geq (|\mathcal{V}|)(|\mathcal{V}| - 1)$ , which implies the lower bound. We now present a formal proof of the lower bound.

## 4.3 Formal Proof of Theorem 4.1

Consider an arbitrary subset  $\mathcal{N} \subset \{1, 2, \dots, N\}$  such that  $|\mathcal{N}| = N - f$ . We construct  $|\mathcal{V}| \times (|\mathcal{V}| - 1)$  executions of the algorithm  $A$ . In particular, for every tuple  $(v_1, v_2) \in \mathcal{V} \times \mathcal{V}$  where  $v_1 \neq v_2$ , we create an execution  $\alpha^{(v_1, v_2)}$  of algorithm  $A$ . In our proof, we first describe execution  $\alpha^{(v_1, v_2)}$  in Section 4.3.1. Then, we describe some properties of execution  $\alpha^{(v_1, v_2)}$  in Section 4.3.2. We use the results of Section 4.3.2 to prove Theorem 4.1 in Section 4.3.3.

### 4.3.1 Execution $\alpha^{(v_1, v_2)}$

In execution  $\alpha^{(v_1, v_2)}$  the readers and the channels from and to the readers do not perform any actions. Among the set of write clients  $\mathcal{C}_w$ , only one write client takes actions. The  $f$  servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  fail at the beginning of the execution. No further server failures occur in the execution. A write  $\pi_1$  is invoked at a write client with value  $v_1$ . All the components of the system except the readers, and the channels from and to the readers, take turns in a fair manner until the completion of  $\pi_1$ . Recall that, in a fair execution where the number of server failures is at most  $f$ , any write that begins eventually terminates irrespective






Figure 2: Pictorial description of executions  $\alpha^{(v_1, v_2)}$  and  $\alpha_i^{(v_1, v_2)}$ .

of the number of read client failures. From the perspective of the servers, write client and the channels between them, the execution  $\alpha^{(v_1, v_2)}$  is indistinguishable from a fair execution where all the read clients fail, we can ensure the execution can be extended until the write operation  $\pi_1$  terminates. Immediately after the termination of  $\pi_1$ , a write operation  $\pi_2$  with value  $v_2$  begins. All the components of the system except the readers and the channels from and to the readers take turns in a fair manner until the completion of  $\pi_2$ . The execution  $\alpha^{(v_1, v_2)}$  ends after the termination of  $\pi_2$ .

#### 4.3.2 Properties of Execution $\alpha^{(v_1, v_2)}$

Let  $P_0^{(v_1, v_2)}, P_1^{(v_1, v_2)}, P_2^{(v_1, v_2)}, \dots, P_M^{(v_1, v_2)}$  be the adjacent points (or points after successive steps) of the constructed execution  $\alpha^{(v_1, v_2)}$ , where  $P_0^{(v_1, v_2)}$  is an arbitrary point after the termination of  $\pi_1$  and before the invocation of  $\pi_2$  and  $P_M^{(v_1, v_2)}$  is an arbitrary point after the point of termination of  $\pi_2$ , and  $M$  is a positive integer. For  $i \in \{0, 1, 2, \dots, M\}$ , we denote by  $\alpha_i^{(v_1, v_2)}$ , the execution between the initial point of  $\alpha^{(v_1, v_2)}$  and point  $P_i^{(v_1, v_2)}$ . The executions  $\alpha^{(v_1, v_2)}$  and  $\alpha_i^{(v_1, v_2)}$  are depicted in Fig. 2.

For an integer  $i$  in  $\{0, 1, \dots, M\}$ , we refer to point  $P_i^{(v_1, v_2)}$  as a  $k$ -valent point if it satisfies certain properties that are described in Definition 4.3,  $k = 1, 2$ . Informally speaking, a point  $P_i^{(v_1, v_2)}$  is said to be  $k$ -valent if there exists an execution that starts at  $P_i^{(v_1, v_2)}$ , where a reader returns  $v_k$ .

**Definition 4.3** ( $k$ -valent,  $k \in \{1, 2\}$ ). For  $i \in \{0, 1, 2, \dots, M\}$ , a point  $P_i^{(v_1, v_2)}$  in the constructed  $\alpha_i^{(v_1, v_2)}$  is said to be  $k$ -valent if we can extend  $\alpha_i^{(v_1, v_2)}$  to an execution  $\beta$  as follows: After  $P_i^{(v_1, v_2)}$  all the messages from and to the writer are delayed indefinitely. A read operation starts at point  $P_i^{(v_1, v_2)}$  and all the components, except the writer and the channels from and to the writer, perform actions until the read operation terminates. The read operation returns  $v_k$ .

It should be noted that a point of an execution can be both 1-valent and 2-valent; thus our definition of valency has a somewhat different structure compared to other definition of valency in other impossibility arguments (e.g. [14]).

**Lemma 4.4.** For  $i \in \{0, 1, 2, \dots, M\}$ , a point  $P_i^{(v_1, v_2)}$  that is not 1-valent is 2-valent.

To show Lemma 4.4, we first prove Lemma 4.5 which informally states that a reader that begins after the termination of the write  $\pi_1$  should return  $v_1$  or  $v_2$  because of regularity of the algorithm.

**Lemma 4.5.** Consider an execution  $\beta$  which is an extension of  $\alpha_i^{(v_1, v_2)}$ . In  $\beta$ , after point  $P_i^{(v_1, v_2)}$ , the writer stops taking steps and all messages from and to the writer are delayed indefinitely. A read operation begins at some point after point  $P_i^{(v_1, v_2)}$  and terminates in  $\beta$ .

Then, the read operation returns either  $v_1$  or  $v_2$ .

The lemma is a natural consequence of the regularity of algorithm  $A$ . We provide a formal proof next.

*Proof.* The read operation is invoked after the termination of write operation  $\pi_1$  in execution  $\beta$ . It is possible that the write operation  $\pi_2$  is invoked before the invocation of the read operation. Because the algorithm is

regular, we must be able to serialize operations in  $\beta$ . Because  $\pi_2$  is invoked after the completion of  $\pi_1$ , the operation  $\pi_2$  is serialized after operation  $\pi_1$ . Regarding the serialization point of the read operation, there are only two possibilities: (i) read operation is serialized immediately after  $\pi_1$  and before  $\pi_2$ , and (ii) the read operation is serialized after  $\pi_2$ . If possibility (i) occurs, that is, if the read operation is serialized after  $\pi_1$ , then it returns  $v_1$ , which is the value of the write operation  $\pi_1$ . If possibility (ii) occurs, then the read returns  $v_2$ , which is the value of the write operation  $\pi_2$ . Therefore, the read operation returns either  $v_1$  or  $v_2$ . This completes the proof.  $\square$

*Proof of Lemma 4.4.* Consider a point  $P_i^{(v_1, v_2)}$  that is not 1-valent. We show that it is 2-valent by constructing an execution  $\beta$  that satisfies Definition 4.3 for  $k = 2$ . The execution  $\beta$  is an extension of  $\alpha_i^{(v_1, v_2)}$ . In  $\beta$ , after point  $P_i^{(v_1, v_2)}$  all the messages from and to the writer are delayed indefinitely.

A read operation  $\pi$  starts at point  $P_i^{(v_1, v_2)}$  in  $\beta$  and all the components, except the writer and the channels from and to the writer, execute their protocols taking turns in a fair manner. From the perspective of the servers, readers and the channels between the servers and the reader, the execution  $\beta$  is indistinguishable from a fair execution of the algorithm where the write client fails before sending or receiving the messages in its channels. Because of the liveness properties satisfied by the algorithm, the read operation terminates. Because the read operation is invoked at point  $P_i^{(v_1, v_2)}$  which is after the point of termination of the write operation  $\pi_1$ , Lemma 4.5 implies that the read operation returns  $v_1$  or  $v_2$ . However, because the point  $P_i^{(v_1, v_2)}$  is not 1-valent, the read operation cannot return  $v_1$ . Therefore the read returns  $v_2$  at some point  $Q$ . Let  $\beta$  denote the extension of  $\alpha_i^{(v_1, v_2)}$  to the point  $Q$ . The execution  $\beta$  satisfies the conditions of Definition 4.3 for  $k = 2$ . Therefore, the point  $P_i^{(v_1, v_2)}$  is 2-valent.  $\square$

**Lemma 4.6.** *There exists some integer  $i \in \{0, 2, \dots, M-1\}$  such that  $P_i^{(v_1, v_2)}$  is 1-valent and  $P_{i+1}^{(v_1, v_2)}$  is not 1-valent.*

*Proof.* To show the lemma, we argue that the following two statements are true.

- (i) Point  $P_0^{(v_1, v_2)}$  is 1-valent.
- (ii) Point  $P_M^{(v_1, v_2)}$  is not 1-valent.

Among all the numbers in  $\{0, 1, 2, \dots, M\}$ , let  $i$  denote the largest number such that  $P_i^{(v_1, v_2)}$  is 1-valent. If (i) is true, we note that the number  $i$  exists. If (ii) is true, then  $i < M$ . Since  $i$  is the largest number such that  $P_i^{(v_1, v_2)}$  is 1-valent, we infer that  $P_i^{(v_1, v_2)}$  is 1-valent, but  $P_{i+1}^{(v_1, v_2)}$  is not 1-valent. The point  $P_i^{(v_1, v_2)}$  therefore satisfies the statement of the lemma. So, to show the statement of the lemma, it suffices to show (i) and (ii). We show (i) and (ii) formally next.

To show (i) we extend  $\alpha_0^{(v_1, v_2)}$  to an execution  $\beta$  as per Definition 4.3 for 1-valency. In  $\beta$ , after point  $P_0^{(v_1, v_2)}$  all the messages from and to the writer are delayed indefinitely.

Note that at point  $P_0^{(v_1, v_2)}$  the write operation  $\pi_2$  has not yet begun. A read operation  $\pi$  starts at point  $P_0^{(v_1, v_2)}$  in execution  $\beta$  and all the components, except the writer and the channels from and to the writer, execute their protocols taking turns in a fair manner. Note that the execution is indistinguishable from a fair execution of the algorithm where the write client fails before sending or receiving the messages in its channels. Because of the liveness properties satisfied by the algorithm, the read operation terminates. Note that there is only one write operation  $\pi_1$  in execution  $\beta$ . Because the algorithm is regular, and because the read operation is invoked after the termination of  $\pi_1$ , it is serialized after  $\pi_1$ . Therefore the read returns  $v_1$  at some point  $Q$ . Let  $\beta$  denote the extension of  $\alpha_0^{(v_1, v_2)}$  to the point  $Q$ . The execution  $\beta$  satisfies the conditions of Definition 4.3 for 1-valency.

To show (ii) we show that we cannot extend  $\alpha_M^{(v_1, v_2)}$  to an execution  $\tilde{\beta}$  as per Definition 4.3. We provide a proof by contradiction. Suppose we can construct  $\tilde{\beta}$  as per Definition 4.3. Note that in  $\tilde{\beta}$ , the read operation is invoked after point  $P_M^{(v_1, v_2)}$ , which is after the point of termination of  $\pi_2$ . Therefore the read operation is invoked after the point of termination of the write  $\pi_2$ . Furthermore,  $\pi_2$  is invoked after the point of termination of  $\pi_1$ . Because  $\tilde{\beta}$  has regular operations, write operation  $\pi_2$  is serialized after write operation  $\pi_1$  and the read is serialized after the write  $\pi_2$ . Therefore the read should return  $v_2$ , which is the value of

write operation  $\pi_2$ . However, because  $\tilde{\beta}$  satisfies Definition 4.3, the read returns  $v_1$  which is not equal to  $v_2$ . Therefore the execution  $\tilde{\beta}$  of algorithm  $A$  does not have regular operations. This is a contradiction to the assumption that the algorithm  $A$  is regular. Therefore point  $P_M^{(v_1, v_2)}$  cannot be 1-valent. This completes the proof.  $\square$

Next, we present the definition of a *pair of critical points* of execution  $\alpha^{(v_1, v_2)}$ .

**Definition 4.7** (Critical points). *Let  $Q_1, Q_2$  be two points in execution  $\alpha^{(v_1, v_2)}$ . The pair of points  $(Q_1, Q_2)$  is defined to be a pair of critical points if there exists a number  $i$  in  $\{0, 2, \dots, M-1\}$  such that*

- $Q_1 = P_i^{(v_1, v_2)}, Q_2 = P_{i+1}^{(v_1, v_2)}$ ,
- $Q_1$  is 1-valent,
- $Q_2$  is not 1-valent.

Lemma 4.6 implies that every execution  $\alpha^{(v_1, v_2)}$  has at least one pair of critical points. Lemma 4.4 implies that if  $(Q_1, Q_2)$  is a pair of critical points in  $\alpha^{(v_1, v_2)}$ , then point  $Q_2$  is 2-valent in  $\alpha^{(v_1, v_2)}$ . We need the following lemma before proceeding to prove Theorem 4.1.

**Lemma 4.8.** *Let  $(Q_1, Q_2)$  be a pair of critical points of execution  $\alpha^{(v_1, v_2)}$ . Then,*

- (a) *the readers, and the channels between the readers and the servers, are all in the same state at point  $Q_2$  as at point  $Q_1$ ;*
- (b) *there is at most one non-failing server  $s$  such that its state at  $Q_1$  is different from its state at  $Q_2$ .*

*Proof.* In execution  $\alpha^{(v_1, v_2)}$ , the readers and the channels between readers and servers do not perform any actions. So these components are in their initial state at every point of the execution, including points  $Q_1$  and  $Q_2$ . This implies that statement (a) is true. We prove (b) next. Note that  $Q_1$  and  $Q_2$  are adjacent points of the execution  $\alpha^{(v_1, v_2)}$ . There are three possibilities: (I) a channel performed an action between points  $Q_1$  and  $Q_2$ , (II) a server performed an action between points  $Q_1$  and  $Q_2$ , or (III) a client performed an action between points  $Q_1$  and  $Q_2$ . We study these three possibilities separately.

**Case I: A channel action took place between points  $Q_1$  and  $Q_2$ .** Note that the algorithm  $A$  does not send any messages on the channels between servers. The channels between readers and servers do not perform any actions in  $\alpha^{(v_1, v_2)}$ . Therefore, we only need to consider the case where a channel between a server and the writer takes an action. If a channel from the writer to server  $s$  takes an action, then, for every server in the set  $\mathcal{N} - \{s\}$ , there was no input, internal or output action between  $Q_1$  and  $Q_2$ . Therefore every server in the set  $\mathcal{N} - \{s\}$ , has the same state at  $Q_1$  as at  $Q_2$ . Similarly, if a channel from a server  $s$  to the writer takes an action, then, for every server in the set  $\mathcal{N} - \{s\}$ , there was no input, internal or output action between  $Q_1$  and  $Q_2$ . Therefore every server in the set  $\mathcal{N} - \{s\}$ , has the same state at  $Q_1$  as at  $Q_2$ . This completes the proof for Case I.

**Case II: A server action took place between points  $Q_1$  and  $Q_2$ .** Let  $s$  be the server that took an action between points  $Q_1$  and  $Q_2$ . This implies that for every server in  $\mathcal{N} - \{s\}$ , no input, output or internal action was taken between these points. Therefore every server in  $\mathcal{N} - \{s\}$ , has the same state at  $Q_1$  as at  $Q_2$ . This completes the proof for Case II.

**Case III: A client action took place between points  $Q_1$  and  $Q_2$ .** If a client action takes place between  $Q_1$  and  $Q_2$ , then, for every server in the system, no input, internal or external action was taken between points  $Q_1$  and  $Q_2$ . Therefore, in this case, every server has the same state at point  $Q_1$  as at point  $Q_2$ . This completes the proof.  $\square$

We are now ready to prove Theorem 4.1.

### 4.3.3 Proof of Theorem 4.1

*Proof of Theorem 4.1.* Lemma 4.6 implies that we can find a pair of critical points  $(Q_1^{(v_1, v_2)}, Q_2^{(v_1, v_2)})$  in execution  $\alpha^{(v_1, v_2)}$ . From Lemma 4.8, we note that there is at most one non-failing server that changes state between  $Q_1^{(v_1, v_2)}$  and  $Q_2^{(v_1, v_2)}$ . Let  $s$  denote the server which changes state between points  $Q_1^{(v_1, v_2)}$  and  $Q_2^{(v_1, v_2)}$ , if there is one; if not, let  $s$  denote an arbitrary non-failing server. For any  $s' \in \mathcal{N}$ , if  $s' \neq s$ , the state of the server  $s'$  is the same at points  $Q_1^{(v_1, v_2)}$  and  $Q_2^{(v_1, v_2)}$ .

Let  $\vec{S}^{(v_1, v_2)}$  be an element of  $\prod_{n \in \mathcal{N}} \mathcal{S}_n \times \mathcal{N} \times \cup_{n \in \mathcal{N}} \mathcal{S}_n$  as follows. The first  $N - f$  components of  $\vec{S}^{(v_1, v_2)}$  denote the states of the  $N - f$  servers in  $\mathcal{N}$  at point  $Q_1^{(v_1, v_2)}$ . The  $(N - f + 1)$ st component of  $\vec{S}^{(v_1, v_2)}$  denote the server index  $s$ , and the  $N - f + 2$ nd component is the state of server  $s$  at point  $Q_2^{(v_1, v_2)}$ . Note that the number of elements in the set  $\bigcup_{(v_1, v_2) \in \mathcal{V} \times \mathcal{V}, v_1 \neq v_2} \{\vec{S}^{(v_1, v_2)}\}$  is at most  $\prod_{i \in \mathcal{N}} |\mathcal{S}_i| \times (N - f) \times \max_{i \in \mathcal{N}} |\mathcal{S}_i|$ .

To prove Theorem 4.1, we show that, if  $(v_1, v_2)$  and  $(v'_1, v'_2)$  are two *distinct* elements of the set  $\{(x, y) : (x, y) \in \mathcal{V} \times \mathcal{V}, x \neq y\}$ , then  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ . If we show this, then it implies that the number of elements in the set  $\bigcup_{(v_1, v_2) \in \mathcal{V} \times \mathcal{V}, v_1 \neq v_2} \{\vec{S}^{(v_1, v_2)}\}$  is at least equal to the number of elements in the set  $\{(x, y) : (x, y) \in \mathcal{V} \times \mathcal{V}, x \neq y\}$ , which is equal to  $(|\mathcal{V}|) \times (|\mathcal{V}| - 1)$ . This leads to the following chain of inequalities:

$$\begin{aligned} \prod_{n \in \mathcal{N}} |\mathcal{S}_n| \times (N - f) \times \max_{n \in \mathcal{N}} |\mathcal{S}_n| &\geq (|\mathcal{V}|) \times (|\mathcal{V}| - 1) \\ \max_{n \in \mathcal{N}} \log_2 |\mathcal{S}_i| + \sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_i| &\geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - \log_2 (N - f) \end{aligned}$$

which implies the theorem. Therefore, to prove the theorem, it suffices to show that, if  $(v_1, v_2) \neq (v'_1, v'_2)$ , then  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ . Suppose, for contradiction, there are two distinct tuples  $(v_1, v_2)$  and  $(v'_1, v'_2)$  in  $\{(x, y) : (x, y) \in \mathcal{V} \times \mathcal{V}, x \neq y\}$  and  $\vec{S}^{(v_1, v_2)} = \vec{S}^{(v'_1, v'_2)}$ .

Let  $i$  denote an integer such that  $Q_1^{(v_1, v_2)}$  is the point  $P_i^{(v_1, v_2)}$  in  $\alpha^{(v_1, v_2)}$ . We now construct executions  $\beta_1^{(v_1, v_2)}$  and  $\beta_2^{(v_1, v_2)}$ , which are extensions of  $\alpha_i^{(v_1, v_2)}$  and  $\alpha_{i+1}^{(v_1, v_2)}$ . Because the point  $Q_1^{(v_1, v_2)}$  is 1-valent, we know there an execution  $\beta_1^{(v_1, v_2)}$  that extends  $\alpha_i^{(v_1, v_2)}$  such that, after point  $Q_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$ , all the messages from and to the writer are delayed indefinitely, and a read operation begins and returns  $v_1$ . Similarly, because the point  $Q_2^{(v_1, v_2)}$  is 2-valent, we know there an execution  $\beta_2^{(v_1, v_2)}$  that extends  $\alpha_{i+1}^{(v_1, v_2)}$  such that, after point  $Q_2^{(v_1, v_2)}$ , all the messages from and to the writer are delayed indefinitely, and a read operation begins and returns  $v_2$ .

The following claim describes a useful property of executions  $\beta_1^{(v_1, v_2)}$  and  $\beta_2^{(v_1, v_2)}$ .

**Claim 4.9.** *Let  $\vec{S}^{(v_1, v_2)} = \vec{S}^{(v'_1, v'_2)}$ . Consider the composite automaton  $\mathcal{A}$  formed by the servers, the readers and the channels between the readers and servers. For  $k \in \{1, 2\}$ , every component of the automaton  $\mathcal{A}$  has the same state at point  $Q_k^{(v_1, v_2)}$  in  $\beta_k^{(v_1, v_2)}$  as at point  $Q_k^{(v'_1, v'_2)}$  in execution  $\beta_k^{(v'_1, v'_2)}$ .*

*Proof of Claim 4.9.* We first consider the case where  $k = 1$ . At points  $Q_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$  and  $Q_1^{(v'_1, v'_2)}$  in  $\beta_1^{(v'_1, v'_2)}$ , all the channels between the readers and servers are empty, the readers are in their initial state and the servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  have failed. Denoting  $\mathcal{N} = \{a_1, a_2, \dots, a_{N-f}\}$  where  $a_1 < a_2 < \dots < a_{N-f}$ , the state of every non-failing server  $a_j, j \in \{1, 2, \dots, N - f\}$  at points  $Q_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$  and  $Q_1^{(v'_1, v'_2)}$  in  $\beta_1^{(v'_1, v'_2)}$  is respectively equal to the  $j$ th component of  $\vec{S}^{(v_1, v_2)}$  and  $\vec{S}^{(v'_1, v'_2)}$ . Because  $\vec{S}^{(v_1, v_2)} = \vec{S}^{(v'_1, v'_2)}$ , every non-failing server is at the same state at  $Q_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$  as at  $Q_1^{(v'_1, v'_2)}$  in  $\beta_1^{(v'_1, v'_2)}$ . This completes the proof for the case where  $k = 1$ .

Consider the case where  $k = 2$ . Let  $s$  denote the server index determined by the  $N - f + 1$ st component of  $\vec{S}^{(v_1, v_2)}$ , which is also equal to the server index determined by the  $N - f + 1$ st component of  $\vec{S}^{(v'_1, v'_2)}$ . All the channels, readers and failed servers have the same state at points  $Q_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  as at  $Q_2^{(v'_1, v'_2)}$ .




Figure 3: Depiction of the proof of Theorem 4.1. (a) Executions  $\beta_k^{(v_1, v_2)}$ ,  $k = 1, 2$ . (b) Executions  $\beta_k^{(v'_1, v'_2)}$ ,  $k = 1, 2$ . (c) Constructed execution  $\gamma$  for the case that  $v'_1 \notin \{v_1, v_2\}$ . (d) Constructed execution  $\gamma$  for the case that  $v'_2 \neq v_2$ .

in  $\beta_2^{(v'_1, v'_2)}$ . The state of every non-failing server except server  $s$  at points  $Q_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  and  $Q_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$  is respectively determined by the corresponding component of  $\vec{S}^{(v_1, v_2)}$  and  $\vec{S}^{(v'_1, v'_2)}$ . The state of server  $s$  at points  $Q_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  and  $Q_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$  are respectively determined by the  $N - f + 2$ nd component of  $\vec{S}^{(v_1, v_2)}$  and  $\vec{S}^{(v'_1, v'_2)}$ . Because  $\vec{S}^{(v_1, v_2)}$  is equal to  $\vec{S}^{(v'_1, v'_2)}$ , every non-failing server is at the same state at  $Q_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  as at  $Q_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$ . This completes the proof of the claim.

*Proof of Theorem 4.1 continued.* We now use Claim 4.9 to obtain a contradiction. Because  $(v_1, v_2)$  and  $(v'_1, v'_2)$  are distinct ordered pairs, there are only two possibilities: (I)  $v'_1 \neq v_1, v'_1 \neq v_2$ , (II)  $v'_2 \neq v_1, v'_2 \neq v_2$ , or  $v'_2 = v_1, v'_1 = v_2$ , both of which imply that  $v'_2 \neq v_2$ . We handle these possibilities separately (See Figure 3).

**Case (I):**  $v'_1 \neq v_1, v'_1 \neq v_2$ .

We create an execution  $\gamma$  of the algorithm  $A$  which contradicts Lemma 4.5. Let  $i$  be an integer such that  $Q_1^{(v_1, v_2)} = P_i^{(v_1, v_2)}$  in execution  $\alpha^{(v_1, v_2)}$ . The execution  $\gamma$  extends execution  $\alpha_i^{(v_1, v_2)}$ , that is, it follows execution  $\alpha^{(v_1, v_2)}$  until point  $Q_1^{(v_1, v_2)}$ . After point  $Q_1^{(v_1, v_2)}$ , the writer, and the channels from and to the writers do not perform any actions. After point  $Q_1^{(v_1, v_2)}$ , the servers, readers and the channels between the servers and readers in  $\gamma$  follow the same steps as the corresponding components in  $\beta_1^{(v'_1, v'_2)}$ .

Claim 4.9 implies that  $\gamma$  is an execution of algorithm  $A$ . In particular,  $\gamma$  is an extension of  $\alpha_i^{(v_1, v_2)}$ , where, after point  $P_i^{(v_1, v_2)}$ , the writer and channels from and to the writer do not perform any actions. From point  $P_i^{(v_1, v_2)}$  onward, since the readers in  $\gamma$  follow the steps of  $\beta_1^{(v'_1, v'_2)}$  until completion, a read begins in  $\gamma$  after this point and terminates returning  $v'_1$ , which is not equal to either  $v_1$  or  $v_2$ . However, according to Lemma 4.5, the read operation in  $\gamma$  should return either  $v_1$  or  $v_2$ . This is a contradiction. Therefore, if  $v_1 \neq v'_1$  and  $v_1 \neq v'_1$ , then  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ .

**Case (II):**  $v'_2 \neq v_2$ .

We create an execution  $\gamma$  of the algorithm  $A$  which leads to a contradiction. Let  $i$  be an integer such that  $Q_1^{(v_1, v_2)} = P_i^{(v_1, v_2)}$  in execution  $\alpha^{(v_1, v_2)}$ . The execution  $\gamma$  extends execution  $\alpha_{i+1}^{(v_1, v_2)}$ , that is, it follows execution  $\alpha^{(v_1, v_2)}$  until point  $Q_2^{(v_1, v_2)}$ . At point  $Q_2^{(v_1, v_2)}$ , the messages from the writers are delayed indefinitely. After point  $Q_2^{(v_1, v_2)}$ , the servers, readers and the channels between the servers and readers in  $\gamma$  follow the same steps as the corresponding components in  $\beta_2^{(v'_1, v'_2)}$ .

Claim 4.9 implies that  $\gamma$  is an execution of algorithm  $A$ . In particular,  $\gamma$  is an extension of  $\alpha_{i+1}^{(v_1, v_2)}$ , where, after point  $P_{i+1}^{(v_1, v_2)}$ , the writer and channels from and to the writer do not perform any actions. From point  $P_{i+1}^{(v_1, v_2)}$  onward, since the readers in  $\gamma$  follow the steps of  $\beta_2^{(v'_1, v'_2)}$  until completion, a read begins in  $\gamma$  after

this point and terminates returning  $v'_2$ , which is not equal to  $v_2$ . However, according to Lemma 4.4 and the fact that  $Q_2^{(v_1, v_2)}$  is not 1-valent,  $Q_2^{(v_1, v_2)}$  is 2-valent, and the read operation in  $\gamma$  should return  $v_2$ . This is a contradiction. Therefore,  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ .

This completes the proof.  $\square$

## 5 A Universal Storage Cost Lower Bound

Our main result of this section is a storage cost lower bound that is applicable to any regular shared memory emulation algorithm, even if it uses server gossip. The lower bound is an implication of Theorem 5.1, which describes constraints on the cardinalities of the server states that must be satisfied by any atomic shared memory emulation algorithm. The lower bounds on the worst case and total storage costs are stated in Corollary 5.2. We provide a sketch of the proof of Theorem 5.1, highlighting the main differences from the proof of Theorem 4.1.

### 5.1 Statement of Theorem 5.1

**Theorem 5.1.** *Let  $A$  be a single-writer-single-reader shared memory emulation algorithm that implements a regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose that every server's state belongs to a set  $\mathcal{S}$  in algorithm  $A$ .*

*Suppose that the algorithm  $A$  satisfies the following liveness property: In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$ , then every operation invoked at a non-failing client terminates. Then,*

$$2 \max_{n \in \mathcal{N}} \log_2 |\mathcal{S}_i| + \sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_i| \geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - 2 \log_2 (N - f)$$

**Corollary 5.2.** *Let  $A$  be a single-writer-single-reader shared memory emulation algorithm that implements a regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose that every server's state belongs to a set  $\mathcal{S}$  in algorithm  $A$ .*

*Suppose that the algorithm  $A$  satisfies the following liveness property: In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$ , then every operation invoked at a non-failing client terminates. Then,*

$$\begin{aligned} \text{MaxStorage}(A) &\geq \frac{\log_2 |\mathcal{V}| + \log_2 |\mathcal{V}| - 1 - 2 \log_2 (N - f)}{N - f + 2}, \\ \text{TotalStorage}(A) &\geq \frac{N(\log_2 |\mathcal{V}| + \log_2 |\mathcal{V}| - 1 - 2 \log_2 (N - f))}{N - f + 2}. \end{aligned}$$

The proof of Corollary 5.2 is similar to the proofs of Corollary 4.2 in Section 4 and Corollary B.2 in Appendix B, and is omitted.

### 5.2 Informal Proof Sketch of Theorem 5.1

The proof of Theorem 5.1 shares many common elements with the proof of Theorem 4.1. The main difference is that now we need to carefully handle the actions performed by the channels between servers. An aspect that distinguishes the proof of Theorem 5.1 from the proof of Theorem 4.1 is that their definitions of  $k$ -valent points. For ease of readability, we inherit the lemmas and definitions from the proof of Theorem 5.1 into this section, so that the proofs can be compared easily. We begin with a proof sketch of Theorem 5.1.

As in our proof of Theorem 4.1, for every subset  $\mathcal{N} \subset \{1, 2, \dots, N\}$  where  $|\mathcal{N}| = N - f$ , we construct an execution  $\alpha^{(v_1, v_2)}$  where the servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  fail at the beginning of the execution. The execution has two write operations for values  $v_1$  and  $v_2$ , where  $v_1 \neq v_2$ . The second write operation with value  $v_2$  begins after the termination of the first write operation with value  $v_1$ .

In this execution, after the point of termination of the first write, if we let the channels between servers deliver all the gossip messages, and then begin a read operation after the delivery of these messages, a reader

can return  $v_1$  because of regularity. Similarly, after the termination of the second write operation, if we let the channels between servers deliver all the gossip messages, a reader can return  $v_2$ . This implies that, in the interval of the second write operation, there are two consecutive points  $P$  and  $P'$  as follows:

- If at point  $P$  we stop the writers and the channels from the writers and let the channels between servers deliver all the gossip messages to arrive at point  $Q$ , then  $v_1$  can be returned by a read operation that begins at point  $Q$ .
- If at point  $P'$  we stop the writers and the channels from the writers and let the channels between servers deliver all the gossip messages to arrive at point  $Q'$ , then  $v_2$  can be returned by a read operation that begins at point  $Q'$ .

By constructing the executions such that gossip messages are delivered in the same order, we can ensure that after the delivery of the messages, at most 2 servers differ in their states between points  $Q$  and  $Q'$ . We use this fact to show that the number of elements in the set of possible server states at points  $Q$  and  $Q'$  points is at most  $\prod_{n \in \mathcal{N}} |\mathcal{S}_i| \times \max_{n \in \mathcal{N}} |\mathcal{S}_n| \times \max_{n \in \mathcal{N}} |\mathcal{S}_n| \times (N - f)^2$ . Therefore, we get  $\prod_{n \in \mathcal{N}} |\mathcal{S}_i| \times (\max_{n \in \mathcal{N}} |\mathcal{S}_n|)^2 \times (N - f)^2 \geq (|\mathcal{V}|)(|\mathcal{V}| - 1)$ , which implies the lower bound.

### 5.3 Formal Proof of Theorem 5.1

#### 5.3.1 Execution $\alpha^{(v_1, v_2)}$ and Its Properties

Let  $\mathcal{N}$  be an arbitrary subset of  $\{1, 2, \dots, N\}$  with  $N - f$  elements. Like the proof of Theorem 4.1, we construct  $|\mathcal{V}| \times (|\mathcal{V}| - 1)$  executions of the algorithm  $A$ . In particular, for every tuple  $(v_1, v_2) \in \mathcal{V} \times \mathcal{V}$  where  $v_1 \neq v_2$ , we create an execution  $\alpha^{(v_1, v_2)}$  of algorithm  $A$ . The execution  $\alpha^{(v_1, v_2)}$  is constructed in a manner that is essentially the same as Section 4.3.1. The  $f$  servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  fail at the beginning of  $\alpha^{(v_1, v_2)}$ . The execution  $\alpha^{(v_1, v_2)}$  has two complete write operations  $\pi_1$  and  $\pi_2$  with values  $v_1$  and  $v_2$ , with  $\pi_2$  being invoked after the termination of  $\pi_1$ .

Similar to the proof of Theorem 4.1, we let  $P_0^{(v_1, v_2)}, P_1^{(v_1, v_2)}, P_2^{(v_1, v_2)}, \dots, P_M^{(v_1, v_2)}$  be a sequence of consecutive points in execution  $\alpha^{(v_1, v_2)}$ , where  $P_0^{(v_1, v_2)}$  is an arbitrary point after the termination of  $\pi_1$  and before the invocation of  $\pi_2$ , and  $P_M^{(v_1, v_2)}$  is an arbitrary point after the point of termination of  $\pi_2$ . We denote by  $\alpha_i^{(v_1, v_2)}$ , the execution between the initial point of  $\alpha^{(v_1, v_2)}$  and point  $P_i^{(v_1, v_2)}$ .

The definition of 1-valent and 2-valent points are similar to Definitions 4.3, with the exception that we allow the channels the servers to deliver their messages before the invocation of the read operation. We provide a formal definition of 1-valent and 2-valent points next.

**Definition 5.3** ( $k$ -valent,  $k \in \{1, 2\}$ ). *For  $i \in \{0, 1, 2, \dots, M\}$ , a point  $P_i^{(v_1, v_2)}$  in the constructed  $\alpha_i^{(v_1, v_2)}$  is said to be  $k$ -valent if we can extend  $\alpha_i^{(v_1, v_2)}$  to an execution  $\beta$  as follows: After  $P_i^{(v_1, v_2)}$  all the messages from and to the writer are delayed indefinitely. At  $P_i^{(v_1, v_2)}$  all the channels between the servers act, delivering all their messages. After the delivery of the messages in the channels between the servers, a read operation starts and all the components, except the writer and the channels from and to the writer, execute their protocols until the read operation terminates. The read operation returns  $v_k$ .*

Results analogous to Lemmas 4.4, 4.5 and 4.6 in the Section 4 hold, with the modified definition of  $k$ -valent points. We simply restate these lemmas without proofs here for the sake of completeness. The proofs are essentially identical to the proofs in Section 4.

**Lemma 5.4.** *For  $i \in \{0, 1, 2, \dots, M\}$ , a point  $P_i^{(v_1, v_2)}$  that is not 1-valent is 2-valent.*

**Lemma 5.5.** *Consider an execution  $\beta$  which is an extension of  $\alpha_i^{(v_1, v_2)}$ . In  $\beta$ , after point  $P_i^{(v_1, v_2)}$ , the writer stops taking steps and all messages from and to the writer are delayed indefinitely. A read operation begins at some point after point  $P_i^{(v_1, v_2)}$  and terminates in  $\beta$ .*

*Then, the read operation returns either  $v_1$  or  $v_2$ .*

**Lemma 5.6.** *There exists some integer  $i \in \{0, 2, \dots, M - 1\}$  such that  $P_i^{(v_1, v_2)}$  is 1-valent and  $P_{i+1}^{(v_1, v_2)}$  is not 1-valent.*

Lemma 4.8 requires some minor modifications to include the possibility that, between a pair of critical points, a channel between servers may perform an action. We state and prove the analogous lemma next.

We inherit the definition of critical points from Definition 4.7. The only change is that the terms 1-valent and 2-valent points use Definition 5.3. We restate the definition here for the sake of completeness.

**Definition 5.7** (Critical points). *Let  $Q_1, Q_2$  be two points in execution  $\alpha^{(v_1, v_2)}$ . The pair of points  $(Q_1, Q_2)$  is defined to be a pair of critical points if there exists a number  $i$  in  $\{0, 2, \dots, M-1\}$  such that*

- $Q_1 = P_i^{(v_1, v_2)}, Q_2 = P_{i+1}^{(v_1, v_2)}$ ,
- $Q_1$  is 1-valent,
- $Q_2$  is not 1-valent.

**Lemma 5.8.** *Let  $(Q_1, Q_2)$  be a pair of critical points of execution  $\alpha^{(v_1, v_2)}$ . Then,*

- (a) *the readers, and the channels between the readers and the servers, are all in the same state at point  $Q_2$  as at point  $Q_1$ ;*
- (b) *there is at most one non-failing server such that its state at  $Q_1$  is different from its state at  $Q_2$ .*
- (c) *among all the channels between the servers, there is at most one channel whose state at  $Q_1$  is different from its state at  $Q_2$ .*

*Proof.* In execution  $\alpha^{(v_1, v_2)}$ , the readers and the channels between readers and servers do not perform any actions. So these components are in their initial state at every point of the execution, including points  $Q_1$  and  $Q_2$ . This implies that statement (a) is true. We prove (b) and (c) next. Note that  $Q_1$  and  $Q_2$  are adjacent points of the execution  $\alpha^{(v_1, v_2)}$ . There are three possibilities: (I) a channel performed an action between points  $Q_1$  and  $Q_2$ , (II) a server performed an action between points  $Q_1$  and  $Q_2$ , or (III) a client performed an action between points  $Q_1$  and  $Q_2$ .

**Case I:** The channels between readers and servers do not perform any actions in  $\alpha^{(v_1, v_2)}$ . We consider two sub-cases: (i) a channel between a server and the writer takes an action, or (ii) a channel between a server and another server takes an action.

*Case I (i):* If a channel from the writer to server  $s$  takes an action, then, for every server in the set  $\mathcal{N} - \{s\}$ , there was no input, internal or output action between  $Q_1$  and  $Q_2$ . Therefore every server in the set  $\mathcal{N} - \{s\}$ , has the same state at  $Q_1$  as at  $Q_2$ . If a channel from a server  $s$  to the writer takes an action, then, for every server, there was no input, internal or output action between  $Q_1$  and  $Q_2$ . Therefore every server in the set  $\mathcal{N} - \{s\}$ , has the same state at  $Q_1$  as at  $Q_2$ . Furthermore, all channels between the servers are in the same state at  $Q_1$  as at  $Q_2$ . This completes the proof in Case I (i).

*Case I (ii):* If a channel from a server, say  $s'$  to a server, say  $s$  takes an action, between  $Q_1$  and  $Q_2$ , then, every server in the set  $\mathcal{N} - \{s\}$ , has the same state at  $Q_1$  as at  $Q_2$ . Furthermore, all the channels between the servers except the channel from  $s'$  to  $s$  have the same state at  $Q_2$  as at  $Q_1$ . Therefore (b) and (c) are satisfied in Case I.

**Case II:** The proof of statement (b) is the same as the proof of Lemma 5.8. We show statement (c) here. Suppose server  $s$  takes an action between  $Q_1$  and  $Q_2$ . If the action is not an output action, or if the server outputs a message onto a channel to the writer, then the states of all the channels between servers are the same at  $Q_1$  and  $Q_2$ . If the action outputs a message on to the channel to another server, say  $s'$ , then all the channels between the servers except the channel from  $s$  to  $s'$  have the same state at  $Q_2$  as at  $Q_1$ . Therefore statement (c) is true for Case II.

**Case III:** The proof of statement (b) is the same as the proof of Lemma 5.8. Statement (c) holds because all the channels between the servers have the same state at  $Q_1$  as at  $Q_2$ .  $\square$

We are now ready to prove Theorem 5.1.



### 5.3.2 Proof of Theorem 5.1

*Proof of Theorem 5.1.* Lemma 5.6 implies that we can find a pair of critical points  $(Q_1^{(v_1, v_2)}, Q_2^{(v_1, v_2)})$  in execution  $\alpha^{(v_1, v_2)}$ . From Lemma 5.8, we note that there is at most one non-failing server and one channel that changes its state between  $Q_1^{(v_1, v_2)}$  and  $Q_2^{(v_1, v_2)}$ . Let  $s$  denote the non-failing server which changes state between points  $Q_1^{(v_1, v_2)}$  and  $Q_2^{(v_1, v_2)}$ , if there is one; if not, let  $s$  denote an arbitrary non-failing server. Let  $s'$  be a non-failing server such that the channel from  $s''$  to  $s'$  change its state between  $Q_1^{(v_1, v_2)}$  and  $Q_2^{(v_1, v_2)}$ , for some server  $s''$ , if there is such a server<sup>6</sup>; let  $s'$  be an arbitrary non-failing server if there is not.

We know that  $Q_1^{(v_1, v_2)}$  is the point  $P_i^{(v_1, v_2)}$  for some  $i \in \{0, 1, \dots, M-1\}$ . We create executions  $\beta_1^{(v_1, v_2)}$  and  $\beta_2^{(v_1, v_2)}$ , which are respectively extensions of  $\alpha_i^{(v_1, v_2)}$  and  $\alpha_{i+1}^{(v_1, v_2)}$  next. The construction of executions  $\beta_1^{(v_1, v_2)}$  and  $\beta_2^{(v_1, v_2)}$  has subtle, but important differences from corresponding constructions in the proof of Theorem 4.1 because, here, we carefully handle the actions of the channels between the servers. After presenting the constructions of  $\beta_1^{(v_1, v_2)}$  and  $\beta_2^{(v_1, v_2)}$ , we state Claim 5.9, which is analogous to Claim 4.9.

Because  $Q_1^{(v_1, v_2)}$  is a 1-valent point, there exists an execution  $\beta_1^{(v_1, v_2)}$  which is an extension of  $\alpha_i^{(v_1, v_2)}$  such that, at point  $P_i^{(v_1, v_2)}$ , the writer and channels from the writer stop performing actions, the channels between the servers deliver all their messages, a read operation begins after the delivery of these messages and returns  $v_1$ . We denote by  $R_1^{(v_1, v_2)}$  a point in  $\beta_1^{(v_1, v_2)}$  after  $P_i^{(v_1, v_2)}$ , after the channels between the servers deliver all their messages, but before the read operation is invoked.

We now create execution  $\beta_2^{(v_1, v_2)}$ . The execution  $\beta_2^{(v_1, v_2)}$  follows  $\alpha^{(v_1, v_2)}$  until point  $Q_2^{(v_1, v_2)}$ . At point  $Q_2^{(v_1, v_2)}$ , all the channels between the servers act delivering all their messages. For a server  $j$  in  $\{1, 2, \dots, N-f\} - \{s, s'\}$ , the channels with destination  $j$  are at the same state at  $Q_2^{(v_1, v_2)}$  as they are at  $Q_1^{(v_1, v_2)}$ ; these channels act and deliver their messages in the same order as they do after point  $Q_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$ . At point  $Q_2^{(v_1, v_2)}$ , server  $j \in \mathcal{N} - \{s, s''\}$  is at the same state as it was at point  $Q_1^{(v_1, v_2)}$ . Also, at point  $Q_2^{(v_1, v_2)}$ , server  $j$  receives messages in the same order as it does at point  $Q_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$ ; on receiving each message, server  $j$  takes the same action in  $\beta_2^{(v_1, v_2)}$  as it does in  $\beta_1^{(v_1, v_2)}$ . The channels with destinations  $s$  or  $s'$  deliver messages in some arbitrary order, and servers  $s$  and  $s'$  perform actions based on the protocol specified by algorithm  $A$ . We denote this point as  $R_2^{(v_1, v_2)}$ . It is worth noting that at point  $R_2^{(v_1, v_2)}$ , all the channels are empty, and every server in  $\mathcal{N} - \{s, s'\}$  is at the same state as it is at point  $R_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$ . After point  $R_2^{(v_1, v_2)}$ , the writer and the channels from and to the writer do not perform any actions. At  $R_2^{(v_1, v_2)}$ , a read operation begins, all the components except the writer and the channels from and to the writer act in a fair manner until the read returns. Because the point  $Q_2^{(v_1, v_2)}$  is 2-valent but not 1-valent, the read returns  $v_2$  in  $\beta_2^{(v_1, v_2)}$ .

We now derive a lower bound on the storage cost by showing some properties on server states at points  $R_1^{(v_1, v_2)}$  and  $R_2^{(v_1, v_2)}$  in executions  $\beta_1^{(v_1, v_2)}$  and  $\beta_2^{(v_1, v_2)}$  respectively.

Let  $\vec{S}^{(v_1, v_2)}$  be an element of  $\prod_{n \in \mathcal{N}} \mathcal{S}_n \times \mathcal{N} \times \cup_{n \in \mathcal{N}} \mathcal{S}_n \times \mathcal{N} \times \cup_{n \in \mathcal{N}} \mathcal{S}_n$  as follows. The first  $N-f$  components of  $\vec{S}^{(v_1, v_2)}$  denote the states of the  $N-f$  servers in  $\mathcal{N}$  at point  $R_1^{(v_1, v_2)}$ . The  $(N-f+1)$ st component of  $\vec{S}^{(v_1, v_2)}$  denotes the server index  $s$  and  $(N-f+2)$ nd component denotes the state of server  $s$  at point  $R_2^{(v_1, v_2)}$  in  $\alpha^{(v_1, v_2)}$ . The  $(N-f+3)$ nd component denotes the server index  $s'$  and the  $(N-f+4)$ th component denotes the state of server  $s'$  at point  $R_2^{(v_1, v_2)}$  in execution  $\beta_2^{(v_1, v_2)}$ . Note that the number of elements in the set  $\bigcup_{(v_1, v_2) \in \mathcal{V} \times \mathcal{V}, v_1 \neq v_2} \{\vec{S}^{(v_1, v_2)}\}$  is at most  $\prod_{i \in \mathcal{N}} |\mathcal{S}_i| \times (N-f)^2 \times (\max_{i \in \mathcal{N}} |\mathcal{S}_i|)^2$ .

To prove Theorem 5.1, we show that, if  $(v_1, v_2)$  and  $(v'_1, v'_2)$  are two *distinct* elements of the set  $\{(x, y) : (x, y) \in \mathcal{V} \times \mathcal{V}, x \neq y\}$ , then  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ . If we show this, then it implies that the number of elements in the set  $\bigcup_{(v_1, v_2) \in \mathcal{V} \times \mathcal{V}, v_1 \neq v_2} \{\vec{S}^{(v_1, v_2)}\}$  is at least equal to the number of elements in the set  $\{(x, y) : (x, y) \in$

<sup>6</sup>Between  $Q_1^{(v_1, v_2)}$ ,  $Q_2^{(v_1, v_2)}$ , if there is a server  $s$  that changes its state, and a channel between two servers  $s''$  and  $s'$  that changes its state, it is easy to show that  $s \in \{s', s''\}$ . We nonetheless use distinct notation for servers  $s, s', s''$  since it simplifies presentation.

$\mathcal{V} \times \mathcal{V}, x \neq y\}$ , which is equal to  $(|\mathcal{V}|) \times (|\mathcal{V}| - 1)$ . This leads to the following chain of inequalities:

$$\begin{aligned} \prod_{n \in \mathcal{N}} |\mathcal{S}_n| \times (N - f)^2 \times (\max_{n \in \mathcal{N}} |\mathcal{S}_n|)^2 &\geq (|\mathcal{V}|) \times (|\mathcal{V}| - 1) \\ 2 \max_{n \in \mathcal{N}} \log_2 |\mathcal{S}_i| + \sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_i| &\geq \log_2 |\mathcal{V}| + \log_2 (|\mathcal{V}| - 1) - 2 \log_2 (N - f) \end{aligned}$$

which implies the theorem. Therefore, to prove the theorem, it suffices to show that, if  $(v_1, v_2) \neq (v'_1, v'_2)$ , then  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ . The remainder of our proof is similar to Theorem 4.1. We highlight the main differences.

Suppose, for contradiction, there are two distinct tuples  $(v_1, v_2)$  and  $(v'_1, v'_2)$  in  $\{(x, y) : (x, y) \in \mathcal{V} \times \mathcal{V}, x \neq y\}$  and  $\vec{S}^{(v_1, v_2)} = \vec{S}^{(v'_1, v'_2)}$ . We now state a lemma that is analogous to Claim 4.9.

**Claim 5.9.** *Let  $\vec{S}^{(v_1, v_2)} = \vec{S}^{(v'_1, v'_2)}$ . Consider the composite automaton formed by the servers, the readers, the channels between the servers, and the channels between the readers and servers. For  $k \in \{1, 2\}$ , every component of this system at point  $R_k^{(v_1, v_2)}$  in  $\beta_k^{(v_1, v_2)}$  is identical to the state of the corresponding component at point  $R_k^{(v'_1, v'_2)}$  in execution  $\beta_k^{(v'_1, v'_2)}$ .*

*Proof of Claim 5.9.* We first consider the case where  $k = 1$ . At points  $R_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$  and  $R_1^{(v'_1, v'_2)}$  in  $\beta_1^{(v'_1, v'_2)}$ , all the channels between the servers, and the channels between the readers and servers are empty, the readers are in their initial state and the servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  have failed. Denoting  $\mathcal{N} = \{a_1, a_2, \dots, a_{N-f}\}$ , The state of every non-failing server  $s$  at points  $R_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$  and  $R_1^{(v'_1, v'_2)}$  in  $\beta_1^{(v'_1, v'_2)}$  is respectively equal to the  $s$ th component of  $\vec{S}^{(v_1, v_2)}$  and  $\vec{S}^{(v'_1, v'_2)}$ . Because  $\vec{S}^{(v_1, v_2)} = \vec{S}^{(v'_1, v'_2)}$ , every non-failing server is at the same state at  $R_1^{(v_1, v_2)}$  in  $\beta_1^{(v_1, v_2)}$  as at  $R_1^{(v'_1, v'_2)}$  in  $\beta_1^{(v'_1, v'_2)}$ . This completes the proof for the case where  $k = 1$ .

Now consider the case where  $k = 2$ . Let  $s$  and  $s'$  respectively denote server indices determined by the  $N - f + 1$ st and  $N - f + 3$ rd components of  $\vec{S}^{(v_1, v_2)}$ . All the channels, readers and failed servers have the same state at points  $R_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  as at  $R_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$ . Denoting  $\mathcal{N} = \{a_1, a_2, \dots, a_{N-f}\}$ , where  $a_1 < a_2 < \dots < a_{N-f}$ , the state of a non-failing server  $a_j \in \{1, 2, \dots, N - f\} - \{s, s'\}$  at points  $Q_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  and  $Q_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$  is respectively equal to the  $j$ th component of  $\vec{S}^{(v_1, v_2)}$  and  $\vec{S}^{(v'_1, v'_2)}$ . The states of servers  $s$  and  $s'$  at point  $Q_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  are respectively determined by the  $N - f + 2$ nd and  $N - f + 4$ th components of  $\vec{S}^{(v_1, v_2)}$ ; the states of  $s$  and  $s'$  at  $Q_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$  are respectively determined by the  $N - f + 2$ nd and  $N - f + 4$ th components of  $\vec{S}^{(v'_1, v'_2)}$ . Because  $\vec{S}^{(v_1, v_2)}$  is equal to  $\vec{S}^{(v'_1, v'_2)}$ , every non-failing server has the same state at  $R_2^{(v_1, v_2)}$  in  $\beta_2^{(v_1, v_2)}$  as at  $R_2^{(v'_1, v'_2)}$  in  $\beta_2^{(v'_1, v'_2)}$ . This completes the proof of the claim.

*Proof of Theorem 5.1 continued.* We now use Claim 5.9 to obtain a contradiction. Because  $(v_1, v_2)$  and  $(v'_1, v'_2)$  are distinct, there are only two possibilities: (I)  $v'_1 \neq v_1, v'_1 \neq v_2$ , (II)  $v'_2 \neq v_1, v'_2 \neq v_2$ , or  $v'_2 = v_1, v'_1 = v_2$ , both of which imply that  $v'_2 \neq v_2$ . We handle these possibilities separately.

**Case (I):**  $v'_1 \neq v_1, v'_1 \neq v_2$ . We create an execution  $\gamma$  as follows. The execution  $\gamma$  follows  $\beta_1^{(v_1, v_2)}$  until point  $R_1^{(v_1, v_2)}$ . From point  $R_1^{(v_1, v_2)}$  in  $\gamma$ , every component except the writer and the channels from and to the writer follows the same steps of the corresponding component in  $\beta_1^{(v'_1, v'_2)}$  from point  $R_1^{(v'_1, v'_2)}$ . Claim 5.9 implies that  $\gamma$  is an execution of algorithm  $A$ . In particular  $\gamma$  extends  $\alpha_i^{(v_1, v_2)}$  such that, the writer stops performing actions, messages from and to the writer are delayed indefinitely and a reader returns  $v'_1$ , which is not equal to  $v_2$  or  $v_1$ . Therefore  $\gamma$  violates Lemma 5.5 and results in a contradiction. Therefore, if  $v'_1 \neq v_1, v'_1 \neq v_2$ , then  $\vec{S}^{(v_1, v_2)} \neq \vec{S}^{(v'_1, v'_2)}$ .

**Case (II):**  $v'_2 \neq v_2$ . We show that the point  $Q_2^{(v_1, v_2)}$  is not 2-valent by constructing an execution  $\gamma$  that satisfies Definition 6.8 for  $k = 2$ . The execution  $\gamma$  follows  $\beta_2^{(v_1, v_2)}$  until point  $R_2^{(v_1, v_2)}$ . From point  $R_2^{(v_1, v_2)}$  in  $\gamma$ , every component except the writer and the channels from and to the writer follows the same steps of the corresponding component in  $\beta_2^{(v'_1, v'_2)}$  from point  $R_2^{(v'_1, v'_2)}$ . Claim 5.9 implies that  $\gamma$  is an execution of

algorithm  $A$ . In particular,  $\gamma$  extends  $\alpha_{i+1}^{(v_1, v_2)}$  such that, after point  $Q_2^{(v_1, v_2)}$ , the messages from and to the writer are delayed indefinitely, the channels between the servers deliver all their messages, and then, a reader begins returning  $v'_2$ , which is not  $v_2$ . However, by Lemma 5.4 and that  $Q_2^{(v_1, v_2)}$  is not 1-valent,  $Q_2^{(v_1, v_2)}$  is 2-valent, and the read of  $\gamma$  should return  $v_2$ , which is a contradiction.

This completes the proof.  $\square$

## 6 Storage Cost Lower Bound for a Restricted Class of Algorithms

In this section, we study a restricted class of algorithms where the write protocols have specific structure. In our restricted class, the write protocols consist of a fixed number of phases. In the protocols that we study, there is only one phase where a message containing information about the actual value is sent to the servers. The formal statement of our assumptions on the write protocol in Section 6.1 is somewhat technically involved. However the write protocols of most previous algorithms [1, 4–6, 11, 12, 21] satisfy our assumptions. After stating our assumptions, we state in Theorem 6.5 in Section 6.2, a storage cost lower bound that applies to the class of algorithms that we study. The lower bound of Theorem 6.5 is much larger than the bound of Theorems 4.1 and 5.1, and is close to the costs of previously developed algorithms.

### 6.1 Protocol Assumptions

We now state three assumptions on the write protocol, Assumptions 1, 2 and 3.

**Assumption 1:** *The state of a write client during a write operation is of the form  $(v, m, h(v, m))$  where  $v \in \mathcal{V}$  is the value of the write operation,  $m$  is an element of a set  $\mathcal{M}$ , and  $h(m, v)$  is the value of function  $h$  whose domain is  $\mathcal{V} \times \mathcal{M}$  and range is a finite set.*

The set  $\mathcal{M}$  is referred to as the metadata set of the write protocol of the algorithm. The function  $h(m, v)$  can contain components of the send buffers that depend on the value, and hashed values used for verification to handle Byzantine adversaries [2, 11, 15]. To describe Assumption 2, we first define the notion of a quorum system and a phase. A quorum system  $\mathcal{Q}$  is a collection of subsets of  $\{1, 2, \dots, N\}$ .

**Definition 6.1** (Phase). *For an arbitrary subset  $\mathcal{N} \subseteq \{1, 2, \dots, N\}$  and a quorum system  $\mathcal{Q}$ , a  $(\mathcal{N}, \mathcal{Q})$ -phase consists of a sequence of actions at a write client as follows: (i) Send message  $m_n$  to server node  $n$  for every  $n \in \mathcal{N}$ . (ii) Wait for responses from least one subset of servers in the collection  $\mathcal{Q}$ . (iii) Perform internal actions, and finish the phase.*

**Definition 6.2** (Decomposable into phases). *A write protocol is said to be decomposable into phases if, on the invocation of a write operation, it invokes a phase, and on the termination of a phase, it either invokes another phase, or terminates the write operation.*

We are now ready to state Assumption 2.

**Assumption 2:** *The write protocol is decomposable into phases.*

Before we state Assumption 3, we state the notion of the black-box action. Informally, a write client action is said to be a black-box action if every internal and output action of the client handles the data values as a black box, that is, actions treat the data object obliviously without regard to the actual value of the object. We state our assumption more formally next. Recall that the write protocol is specified as a set of transitions (old-state, action, new-state).

**Definition 6.3** (Black-box Action). *An internal or output action  $\sigma$  performed by a write client is said to be a black-box action if the following holds: if, for some value  $v \in \mathcal{V}$ ,*

- *the action  $\sigma$  is enabled when the client's state is  $(m, v, h(m, v))$  for some  $m \in \mathcal{M}$ , and*
- *the action  $\sigma$  can result in the transition of the client's state from  $(m, v, h(m, v))$  to  $(m', v, h(m, v))$  for some  $m' \in \mathcal{M}$ ,*

*then, for every value  $v' \in \mathcal{V}$ ,*

- *the action  $\sigma$  is enabled when the client's state is  $(m, v', h(m, v'))$ , and*
- *the action  $\sigma$  can result in the transition of the client's state from  $(m, v', h(m, v'))$  to  $(m', v', h(m', v'))$ .*

For example, in the ABD algorithm [3], all actions are black-box actions. In particular, if the action of sending of a value is enabled when the metadata is  $m$  for particular value  $v$ , then the send action is enabled for every value  $v' \in \mathcal{V}$ . Similarly, in erasure coding based algorithms, if the action of sending a codeword symbol is enabled in at one state, it is enabled at every state.

Note that a write client's output actions are send and return. Send actions of a write client are categorized as value-dependent and value-independent actions.

**Definition 6.4** (Value-dependent and value-independent send actions). *A black-box send action  $\sigma$  that is enabled during a write operation is said to be value-independent if the message sent does not depend on the value of the operation. A send action that is not a value-independent send action is referred to as a value-dependent send action.*

For example, in the ABD algorithm, value-independent send actions involve sending query messages to the servers. Messages sent by value-dependent and value-independent send actions are respectively referred to as value-dependent and value-independent messages. We are now ready to state Assumption 3.

**Assumption 3:** (a) *All write client actions are black-box actions,* and (b) *in a write operation  $\pi$  in an execution  $\alpha$ , if there is a phase where at least one value-dependent send action is performed, then every send action in every subsequent phase of the write operation  $\pi$  is a value-independent send action.*

In particular, Assumption 3(b) implies that there is at most one phase where the writer sends value-dependent messages on behalf of a write operation in any execution. We next state our main result.

## 6.2 Statement of Theorem 6.5

We state our theorem for *weakly regular* MWMR registers [22]. Informally a weakly regular shared memory object is one that supports concurrent write and read operations where, in every execution, for every terminating read operation  $\pi_r$ , there is a subset  $\Phi$  of the non-terminating write operations such that the operations in  $\{\pi_r\} \cup \Phi \cup \Pi$  look like the execution of a serial variable, where  $\Pi$  is the set of all terminating write operations in the execution.

An atomic register is also a weakly regular register. Therefore, the storage cost of Theorem 6.5 applies for atomic registers as well.

In an execution  $\alpha$ , a write operation  $\pi$  is said to be active at point  $P$  if the point  $P$  is after the point of invocation but before the point of termination of  $\pi$ .

**Theorem 6.5.** *Let  $A$  be a multi-writer-single-reader shared memory emulation algorithm that implements a weakly regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose algorithm  $A$  satisfies Assumptions 1, 2 and 3 stated in Section 6.1, and following liveness property: In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$  and the number of active write operations is no bigger than  $\nu$ , then every operation invoked at a non-failing client terminates.*

*Then, for every subset  $\mathcal{N} \subseteq \{1, 2, \dots, N\}$ ,  $|\mathcal{N}| = \min(N - f + \nu - 1, N)$*

$$\sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_n| \geq \log_2 \binom{|\mathcal{V}| - 1}{\nu^*} - \nu^* \log_2 (N - f + \nu^* - 1) - \log_2 (\nu^*!)$$

where  $\nu^* = \min(\nu, f + 1)$ .

**Corollary 6.6.** *Let  $A$  be a multi-writer-single-reader shared memory emulation algorithm that implements a weakly regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose algorithm  $A$  satisfies Assumptions 1, 2 and 3 stated in Section 6.1, and following liveness property: In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$  and the number of write operations is no bigger than  $\nu$ , then every operation invoked at a non-failing client terminates. Then*

$$\begin{aligned} \text{MaxStorage}(A) &\geq \frac{\nu^*}{N - f + \nu^* - 1} \log |\mathcal{V}| - o(\log |\mathcal{V}|), \\ \text{TotalStorage}(A) &\geq \frac{\nu^* N}{N - f + \nu^* - 1} \log |\mathcal{V}| - o(\log |\mathcal{V}|), \end{aligned}$$

where  $\nu^* = \min(\nu, f + 1)$

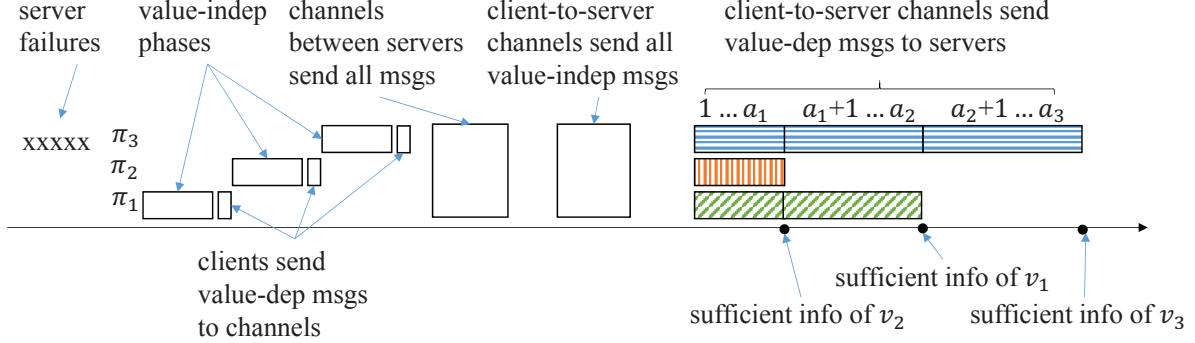


Figure 4: Pictorial description of the execution for  $\nu = 3$ .  $b_1 = 2, b_2 = 1, b_3 = 3$ .

The proof of Corollary 6.6 is similar to the proofs of Corollary 4.2 in Section 4 and Corollary B.2 in Appendix B, and is omitted.

### 6.3 Informal, Intuitive Proof Sketch for Theorem 6.5

For simplicity of exposition, assume  $\mathcal{N} = \{1, 2, \dots, N - f + 2\}$ . For our informal description, we set the parameter  $\nu = 3$ , and  $f \geq \nu - 1 = 2$ . Our proof of Theorem 6.5 constructs an execution  $\alpha$  where the servers in  $\{N - f + 3, N - f + 4, \dots, N\}$  fail at the beginning of the execution. The execution has  $\nu = 3$  write operations  $\pi_1, \pi_2, \pi_3$  with distinct values  $v_1, v_2, v_3$  respectively invoked at distinct clients  $C_1, C_2, C_3$ . We assume that at the beginning of the execution before the invocation of any write operation, a default initial value  $v_0$  can be returned by any read operation, and that values  $v_1, v_2, v_3$  are distinct from the default initial value  $v_0$ .

Writes  $\pi_1, \pi_2, \pi_3$  are invoked respectively at clients  $C_1, C_2, C_3$ . Recall that, as per Assumption 3, there is at most one phase where the clients send value-dependent messages. Operations  $\pi_1, \pi_2, \pi_3$  execute their protocols in a fair manner until they reach their respective phases where they send the value-dependent messages. The clients send the value-dependent messages onto the channels, but the channels do not yet deliver these value-dependent messages. Consider the point  $P$  after all three clients send their value-dependent messages. At point  $P$ , the channels from the clients to the servers carry all the value-dependent messages that can be sent in the execution  $\alpha$ , and none of them are delivered to any of the servers.

Now we construct an execution  $\alpha'$  which extends the execution  $\alpha$  beyond point  $P$  to a point  $P'$  by allowing the channels from the clients to the servers act to deliver all the value-dependent messages to the first  $N - f$  servers at point  $P$ . After the delivery of the messages, it must be the case that the first  $N - f$  servers store “sufficient information” to return at least one of the values  $v_1, v_2$  or  $v_3$ . This is because there are no additional phases where value-dependent messages are sent in  $\alpha'$ , and so, even if we extend the execution beyond point  $P'$ , the servers cannot receive any additional information related to  $v_1, v_2$  or  $v_3$ . Furthermore, if we extend the execution  $\alpha'$  beyond  $P'$  by letting at least one of the operations  $\pi_1, \pi_2, \pi_3$  complete by performing the remaining phases, then, because of weak regularity, at least one of the values  $v_1, v_2$  or  $v_3$  must be returnable from the first  $N - f$  servers after the completion of the operation. So it must be the case that at point  $P'$  in  $\alpha'$ , the servers store sufficient information to return one of  $v_1, v_2$  or  $v_3$ . Let  $a_1$  be the smallest number such that, if the channels between the clients and the first  $a_1$  servers act after point  $P$  by delivering all their messages, then the first  $a_1$  servers store sufficient information of value  $v_{b_1}$  for some  $b_1 \in \{1, 2, 3\}$ . Note that  $1 \leq a_1 \leq N - f$ . In our execution  $\alpha$ , at point  $P$ , we let all the channels deliver all their value-dependent messages to the first  $a_1$  servers. Denote the point after the delivery of the messages as  $P_1$ . Since we chose  $a_1$  to be the smallest number of servers that contain sufficient information of any one of  $v_1, v_2, v_3$ , sufficient information of any *one* of  $v_1, v_2$  or  $v_3$  is not contained from any of the first  $a_1 - 1$  servers at point  $P_1$ .

Now we construct an execution  $\alpha''$  as an extension of execution  $\alpha$  beyond  $P_1$  by allowing the channels from clients  $\{C_1, C_2, C_3\} - \{C_{b_1}\}$  deliver their value-dependent messages to servers in  $\{a_1 + 1, a_1 + 2, \dots, N - f + 1\}$ . After the delivery of the messages, sufficient information of one of the values  $v_{b_2}$  must be available in the

first  $N - f + 1$  servers for some  $b_2 \in \{1, 2, 3\} - \{b_1\}$ . This is because, if, after the delivery of the value-dependent messages in  $\alpha''$ , server  $a_1$  stops taking actions, and clients  $\{C_1, C_2, C_3\} - \{C_{b_1}\}$ , and the first  $N - f + 1$  servers apart from server  $a_1$  take actions in a fair manner, then one of the operations  $\pi_1, \pi_2, \pi_3$  completes; weak regularity implies that one of the values  $v_1, v_2, v_3$  is returnable from the first  $N - f + 1$  servers. However, note that sufficient information related to  $v_{b_1}$  is not contained in the first  $a_1 - 1$  servers. As a consequence,  $v_{b_1}$  cannot be returnable from the first  $N - f + 1$  servers in  $\alpha''$  if server  $a_1$  does not take actions and we do not allow the value-dependent messages from client  $C_{b_1}$  to be delivered to any one of the servers  $\{a_1 + 1, a_1 + 2, \dots, N - f + 1\}$ ; therefore a value  $v_{b_2} \neq v_{b_1}$  must be returnable from the first  $N - f + 1$  servers. Let  $a_2$  be a number with  $a_1 < a_2 \leq N - f + 1$  such that, if all the channels deliver their value-dependent messages to the first  $a_1$  servers and the channels from clients in  $\{C_1, C_2, C_3\} - \{C_{b_1}\}$  deliver their value-dependent messages to the servers in  $\{a_1 + 1, a_1 + 2, \dots, a_2\}$ , then sufficient information about value  $v_{b_2}$  is contained in the first  $a_2$  servers for some  $b_2 \neq b_1, b_2 \in \{1, 2, 3\}$ . In  $\alpha$ , at point  $P_1$ , we let clients in  $\{C_1, C_2, C_3\} - \{C_{b_1}\}$  deliver their value-dependent messages to the servers in  $\{a_1 + 1, a_1 + 2, \dots, a_2\}$ . Denote the point after the delivery of the messages as  $P_2$ .

Similarly, if we let the channels from remaining client in  $\{C_1, C_2, C_3\} - \{C_{b_1}, C_{b_2}\}$  deliver their value-dependent messages at point  $P_2$  in  $\alpha$  to the servers in  $\{a_2 + 1, a_2 + 2, \dots, N - f + 2\}$ , then sufficient information about the value in  $\{v_1, v_2, v_3\} - \{v_{b_1}, v_{b_2}\}$  is contained from the first  $N - f + 2$  servers after the delivery of the messages. At this point, sufficient information about all 3 values is contained in the first  $N - f + 2$  servers. We can show that this implies that there is a one-to-one mapping from the states of the first  $N - f + 2$  servers to the values in  $(\mathcal{V} - \{v_0\})^3$ , where  $v_0$  is the initial value. This implies that the storage cost must be at least  $\frac{3}{N-f+2} \log_2 |\mathcal{V}| + o(\log_2 |\mathcal{V}|)$ .

Our proof involves developing an appropriate notion of *sufficient information of a value* that is applicable even when each server stores some arbitrary function of the values of the different versions it receives. In particular, we cannot directly borrow from other work [23], whose notion of sufficient information of a value is tied to the storage scheme imposed by the model studied. Our notion of sufficient information is crystallized in a notion of valency that is more general as compared with Section 4.

## 6.4 Proof

To avoid cumbersome notation, we assume that  $\nu \leq f + 1$  and we prove Theorem 6.5 for  $\mathcal{N} = \{1, 2, \dots, N - f + \nu - 1\}$ . The proof for the general case readily follows from the proof provided here.

To prove the lower bound, we construct a set of executions as follows. Every element of the set is parametrized by:

- an arbitrary permutation  $\sigma : \{1, 2, \dots, \nu\} \rightarrow \{1, 2, \dots, \nu\}$ ,
- an arbitrary collection of numbers  $a_1, a_2, \dots, a_\nu \in \{0, 1, \dots, N - f + \nu - 1\}$ , where  $a_1 \leq a_2 \leq \dots \leq a_\nu$ , and
- an arbitrary collection of distinct values  $v_0, v_1, v_2, \dots, v_\nu \in \mathcal{V}$ .

An execution parametrized by permutation  $\sigma$ , numbers  $a_1, \dots, a_\nu$  and values  $v_0, v_1, \dots, v_\nu$  is denoted by  $\alpha^{(v_0, v_1, \dots, v_\nu)}(\sigma, a_1, \dots, a_\nu)$ . We assume that  $v_0$  indicates the default initial value that should be returned by a read operation in an execution where there is no write operation.

The proof is split into four parts. In the first part, we describe our construction of execution  $\alpha^{(v_0, v_1, \dots, v_\nu)}(\sigma, a_1, \dots, a_\nu)$  in Section 6.4.1. We also prove a property, Lemma 6.9, regarding the states of the components in execution  $\alpha^{(v_0, v_1, \dots, v_\nu)}(\sigma, a_1, \dots, a_\nu)$ . In the second part, we define the notion of valency tailored to our class of executions in Section 6.4.2. In the third part, we prove a key property in Lemma 6.10 in Section 6.4.3. In the fourth and final part, we use Lemma 6.10 to prove Theorem 6.5.

**Notation:** In the sequel, we denote  $\vec{v} = (v_0, v_1, \dots, v_\nu)$ , where  $v_0, v_1, v_2, \dots, v_\nu \in \mathcal{V}$ . We use the term *value-vector* to refer to the  $\vec{v}$ . We assume that the components of a value vector are distinct, and  $v_0$  is the default initial value.

#### 6.4.1 Description of Execution $\alpha^{\vec{v}}(\sigma, a_1, \dots, a_\nu)$

We next describe the execution  $\alpha^{\vec{v}}(\sigma, a_1, \dots, a_\nu)$  for an arbitrary value vector  $\vec{v}$ . In the execution, only  $\nu$  distinct write clients  $C_1, \dots, C_\nu \in \mathcal{C}_w$  act. In particular, for  $i \in \{1, 2, \dots, \nu\}$ , one write operation  $\pi_i$  is invoked at client  $C_i$  with value  $v_i$ . For every collection of integers  $a_1, a_2, \dots, a_\nu \in \{0, 1, \dots, N - f + \nu - 1\}$ , where  $a_1 \leq a_2 \leq \dots \leq a_\nu$ , and for every permutation  $\sigma$ , the execution  $\alpha^{\vec{v}}(\sigma, a_1, \dots, a_\nu)$  is an extension of an execution  $\alpha_0^{\vec{v}}$ . The final point of  $\alpha_0^{\vec{v}}$  is denoted as  $P_0^{\vec{v}}$ . We first describe  $\alpha_0^{\vec{v}}$  and later describe  $\alpha^{\vec{v}}(\sigma, a_1, \dots, a_\nu)$ .

We choose arbitrarily a reference value vector  $\vec{v}_{\text{ref}} \in \{v_0\} \times \mathcal{V}^\nu$ . The components of  $\vec{v}_{\text{ref}}$  are denoted as  $v_0, v_{1,\text{ref}}, v_{2,\text{ref}}, \dots, v_{\nu,\text{ref}}$ . Next, we construct execution  $\alpha_0^{\vec{v}_{\text{ref}}}$ . After that, we use our construction of  $\alpha_0^{\vec{v}_{\text{ref}}}$  to describe the execution  $\alpha_0^{\vec{v}}$  for an arbitrary value vector  $\vec{v}$ .

---

Execution  $\alpha_0^{\vec{v}_{\text{ref}}}$

---

- 1: Initial point: all components are at their initial states.
  - 2: The last  $f + 1 - \nu$  servers fail
  - 3: **for**  $i = 1$  to  $\nu$  **do**
  - 4:     Operation  $\pi_i$  is invoked at client  $C_i$  with value  $v_{i,\text{ref}}$ .
  - 5:     Client  $C_i$ , the non-failed servers and the channels take steps in a fair manner until the beginning of a phase  $R_i$ , where at least one value-dependent send action is enabled, or until operation  $\pi_i$  terminates without sending any value-dependent message.
  - 6:     If operation  $\pi_i$  is not terminated, client  $C_i$  performs the send actions corresponding to phase  $R_i$ , sending all value-dependent message to the channels. (The channels from the client to the servers do not yet deliver the value-dependent messages. The client  $C_i$  does not perform any more actions.)
  - 7: **end for**
  - 8: The channels between the servers deliver all their messages.
  - 9: The channels from the clients to the servers deliver all the value-independent messages.
- 

Note that in execution  $\alpha_0^{\vec{v}_{\text{ref}}}$ , until the beginning of phase  $R_i$  at client  $C_i$  for any  $i \in \{1, 2, \dots, \nu\}$ , every action is a value-independent action. The only value-dependent send actions in the execution are the send actions performed by clients  $C_i$  in their corresponding phases  $R_i$ , for  $i \in \{1, 2, \dots, \nu\}$ . Note that these value-dependent messages are not delivered in the execution  $\alpha_0^{\vec{v}}$ . Therefore, from the perspective of the servers, all the received messages by the final point  $P_0^{\vec{v}_{\text{ref}}}$  are value-independent messages.

We now describe execution  $\alpha_0^{\vec{v}}$  for an arbitrary value vector  $\vec{v}$ . In execution  $\alpha_0^{\vec{v}}$ , the behavior of the environment, servers, clients, and is the same as in execution  $\alpha$ . That is, for every point  $P^{\text{ref}}$  in  $\alpha_0^{\vec{v}_{\text{ref}}}$ , there is a corresponding point  $P$  in  $\alpha_0^{\vec{v}}$ . If at  $P^{\text{ref}}$  there is a write invoked at client  $C_i$  with value  $v_{i,\text{ref}}$  for some  $i \in \{1, 2, \dots, \nu\}$ , then at point  $P$  in  $\alpha_0^{\vec{v}}$ , there is write invoked at client  $C_i$  with value  $v_i$ . If at  $P^{\text{ref}}$  a channel or a server performs an action, then at point  $P$ , the same channel or server performs the same action. If at  $P^{\text{ref}}$  a client performs a black-box action  $\sigma$ , then at point  $P$ , the same client performs the action  $\sigma$  such that it takes the corresponding internal transition as in Definition 6.3. We next argue that  $\alpha_0^{\vec{v}}$  is a valid execution of the algorithm. Since read clients do not act in execution  $\alpha^{\vec{v}}$ , we only argue that servers and write clients conform to their protocol specifications in the execution.

**Lemma 6.7.** *The execution  $\alpha_0^{\vec{v}}$  is a valid execution of the algorithm for every value vector  $\vec{v}$ .*

*Proof.* The servers, channels and clients in the system are I/O automata. Let  $\mathcal{A}_s$  denote the automaton formed from the composition of all the servers and the channels between servers. Let  $\mathcal{A}_c$  denote the automaton formed by the composition of all the write clients.

The input to the automaton  $\mathcal{A}_s$  are the messages delivered in the channels from the clients to the servers. Since only value-independent messages are delivered to the servers in executions  $\alpha_0^{\vec{v}_{\text{ref}}}$  and  $\alpha_0^{\vec{v}}$ , the inputs to  $\mathcal{A}_s$  in the two executions are the same. Since the components of  $\mathcal{A}_s$  follow their protocol specifications in execution  $\alpha_0^{\vec{v}_{\text{ref}}}$ , and the steps of the components in  $\alpha_0^{\vec{v}}$  are the same as in execution  $\alpha_0^{\vec{v}_{\text{ref}}}$  the components of  $\mathcal{A}_s$  follow their protocol specification in execution  $\alpha_0^{\vec{v}}$  as well.

The inputs to components of  $\mathcal{A}_c$  are the messages sent from the servers to the write clients, and the write invocations. Because the write invocations and the server actions in execution  $\alpha_0^{\vec{v}}$ , are the same as in execution  $\alpha_0^{\vec{v}_{\text{ref}}}$ , the inputs to the automaton  $\mathcal{A}_c$  are the same. Since all write client actions are black-box actions, and since write client steps follow their protocol specifications in  $\alpha_0^{\vec{v}_{\text{ref}}}$ , the clients follow their protocol specifications in execution  $\alpha_0^{\vec{v}}$  as well. Therefore, execution  $\alpha_0^{\vec{v}}$  is a valid execution of the algorithm.  $\square$

It is useful to note that, at the final point  $P_0^{\vec{v}}$  of  $\alpha_0^{\vec{v}}$ , the states of the servers, the channels amongst the servers, the channels from the servers to the clients, and the metadata components of the client states are all independent of the value vector. The only components whose state may depend on the value vector are the write clients, and the channels from the write clients to the servers. The following lemma describes this property more formally.

**Lemma 6.8.** *For any two value vectors  $\vec{v}, \vec{v}'$ , the state of every server, every channel from a server to a server or client, and the metadata components of the state of any writer  $C \in \{C_1, C_2, \dots, C_\nu\}$  at the final point  $P_0^{\vec{v}}$  of execution  $\alpha_0^{\vec{v}}$  is the same as at the final point  $P_0^{\vec{v}'}$  of  $\alpha_0^{\vec{v}'}$ .*

*Proof.* Consider a component automaton  $c$  the system, which is a server, or a channel from a servers to a server or a client. We consider four possibilities separately.

*c is channel from a server to another server:* At point  $P_0^{\vec{v}}$  in  $\alpha_0^{\vec{v}}$ , as at point  $P_0^{\vec{v}'}$  in  $\alpha_0^{\vec{v}'}$  the channel  $c$  is empty. Therefore, for a component that is a channel between two servers, the lemma statement holds.

*c is a server:* Server  $c$  takes the same steps in  $\alpha_0^{\vec{v}}$ , as in  $\alpha_0^{\vec{v}'}$ . Therefore the state of  $c$  at point  $P_0^{\vec{v}}$  in  $\alpha_0^{\vec{v}}$  is the same as its state  $P_0^{\vec{v}'}$  in  $\alpha_0^{\vec{v}'}$ .

*c is channel from a server to a client:* If  $c$  is a channel from a server  $s$  to a client, note that the server  $s$  takes the same steps in  $\alpha_0^{\vec{v}}$ , as in  $\alpha_0^{\vec{v}'}$ . Therefore, the inputs to channel  $c$  in the two executions are the same. The steps of the channel are the same in the two executions as well. Therefore, for a component that is a channel from a server to a client, the lemma statement holds.

*c is a client  $C_i$  for some  $i \in \{1, 2, \dots, \nu\}$ :* The client  $C_i$  takes the same steps in  $\alpha_0^{\vec{v}}$  as in  $\alpha_0^{\vec{v}'}$ , except for its final action in the execution. The final action of the execution is either a value-dependent send action, or an operation termination action. In either case, because all actions of client  $C_i$  are black-box actions, and the client  $C_i$  performs the same action that is performed in  $\alpha_0^{\vec{v}_{\text{ref}}}$ , the metadata component of the client  $C_i$  after the final action is the same at point  $P_0^{\vec{v}}$  in  $\alpha_0^{\vec{v}}$ , as at point  $P_0^{\vec{v}'}$  in  $\alpha_0^{\vec{v}'}$ .  $\square$

We describe  $\alpha^{\vec{v}}(\sigma, a_1, \dots, a_\nu)$  as an extension of  $\alpha_0^{\vec{v}}$ . If  $a_1 \geq 1$ , then at point  $P_0^{\vec{v}}$  of execution  $\alpha^{\vec{v}}(\sigma, a_1, \dots, a_\nu)$ , for every server node  $n$  in  $\{1, 2, \dots, a_1\}$ , the channels from all the writers to server  $n$  deliver their messages. We denote the point after the delivery of all messages as  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ . After the delivery of the messages, the following actions take place:

**for**  $i = 1$  to  $i = \nu - 1$  **do**

**if**  $a_{i+1} > a_i$  **then** For every server node  $n$  in  $\{a_i + 1, a_i + 2, \dots, a_{i+1}\}$  the channels from every writer in  $\{C_1, C_2, \dots, C_\nu\} - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\}$  to server  $n$  deliver all their messages. We denote the point after the delivery of all the messages as  $P_{i+1}^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ .

**end if**

**end for**

Based on the above procedure, note that a server  $n$  that belongs to  $\{a_i + 1, a_i + 2, \dots, a_{i+1}\}$  receives value-dependent messages from all clients except the clients in  $C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}$ . For  $i \in \{0, 1, 2, \dots, \nu\}$ , the portion of the execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  from the initial point until point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is denoted as  $\alpha_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ .

The following property is due to Lemma 6.8 and the fact that the only client-to-server channels act after point  $P_0^{\vec{v}}$ .

**Lemma 6.9.** *For two value vectors  $\vec{v}, \vec{v}'$ , for any two permutations  $\sigma, \bar{\sigma}$  and integers  $0 \leq a_1 \leq a_2 \leq \dots a_\nu \leq N - f + \nu - 1$  and  $0 \leq \bar{a}_1 \leq \bar{a}_2 \leq \dots \bar{a}_\nu \leq N - f + \nu - 1$  and integers  $0 \leq i_0, j_0 \leq \nu$ , the state of every*



channel from a server to a server or a client, and the metadata components of the state of every writer in  $\{C_1, C_2, \dots, C_\nu\}$  is the same at point

$$P_{i_0}^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$$

in execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , as at point

$$P_{j_0}^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$$

in execution  $\alpha^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$ .

*Proof.* Consider a component  $c$  which is a channel between two servers, or a channel from a server to a client, or a write client. The component  $c$  has the same state at point

$$P_{i_0}^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$$

in execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  as at point  $P_0^{\vec{v}}$  of the execution, since it does not take any steps between  $P_0^{\vec{v}}$  and  $P_{i_0}^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ . Similarly, the component  $c$  has the same state at point

$$P_{j_0}^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$$

as at point  $P_0^{\vec{v}'}$  in execution  $\alpha^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$ .

Lemma 6.8 implies that if  $c$  is a channel between servers, or a channel from a server to a client, then it has the same state at point  $P_0^{\vec{v}}$  in execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  as at point  $P_0^{\vec{v}'}$  in execution  $\alpha^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$ . This implies that  $c$  has the same state at point

$$P_{i_0}^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$$

in execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , as at point

$$P_{j_0}^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$$

in execution  $\alpha^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$ .

Similarly, Lemma 6.8 implies that if  $c$  is a client, then its metadata component is the same at point

$$P_{i_0}^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$$

in execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , as at point

$$P_{j_0}^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$$

in execution  $\alpha^{\vec{v}'}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu)$ . This completes the proof.  $\square$

#### 6.4.2 Definition of Valency

Consider a collection of distinct values  $v_0, v_1, \dots, v_\nu \in \mathcal{V}$ , a permutation  $\sigma$  and a collection of numbers  $a_1, a_2, \dots, a_\nu \in \{0, 1, 2, \dots, N - f + \nu - 1\}$ , where  $0 \leq a_1 \leq a_2 \leq \dots \leq a_\nu \leq N - f + \nu - 1$ . For a subset of write clients  $\mathcal{C}_0 \subseteq \{C_1, C_2, \dots, C_\nu\}$ , and  $1 \leq j \leq \nu$ , the point  $P$  in execution  $\alpha^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is said to be  $(j, \mathcal{C}_0)$ -valent if there exists some execution  $\beta$  which is an extension of  $\alpha_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  such that, after  $P$  in  $\beta$ ,

- the writers in  $\mathcal{C}_w - \mathcal{C}_0$  do not send any value-dependent messages, the channels from the writers in  $\mathcal{C}_w - \mathcal{C}_0$  do not deliver any value-dependent messages, and
- there is a read operation that begins at a reader and completes returning value  $v_j$ .

A  $(j, \{\})$ -valent point is simply referred to as a  $j$ -valent point. It is instructive to note that a point that is  $j$ -valent is also  $(j, \mathcal{C}_0)$ -valent for any subset  $\mathcal{C}_0 \subseteq \{C_1, C_2, \dots, C_\nu\}$ .

*Intuition for the definition of valency:* Before proceeding, we provide an intuitive explanation of the notion of valency. Consider a point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  which is  $(1, C_2)$ -valent. Intuitively, at such a point, the servers already have “sufficient information” to retrieve value  $v_1$ . However, to recover  $v_1$ , value-dependent actions from client  $C_2$  or the channels from  $C_2$  could be necessary. To see this, consider the following scenario: in the execution some algorithm, at point  $P$ , every server stores  $v_1 + v_2$ , where the set  $\mathcal{V}$  is interpreted to be some finite field, and  $+$  indicates the addition operator over the field. In this case, in general, neither value  $v_1$  nor value  $v_2$  is retrievable from the system. However, *given value*  $v_2$ , it can be subtracted off and the value  $v_1$  would be retrievable. A clever protocol could ensure that the value-dependent messages of client  $C_2$  will be used to subtract  $v_2$  from a sufficient number of servers to ensure that  $v_1$  is returnable, even if client  $C_1$  did not take any value-dependent actions. In this case the point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  would be considered  $(1, C_2)$ -valent.

### 6.4.3 A Key Lemma

The following lemma is a key component of our proof of Theorem 6.5.

**Lemma 6.10.** *Let  $\prec$  be a total ordering on  $\mathcal{V}$ . Given a collection of distinct values  $v_1, v_2, \dots, v_\nu \in \mathcal{V}$ , there is a permutation  $\sigma : \{1, 2, \dots, \nu\} \rightarrow \{1, 2, \dots, \nu\}$ , and a collection of distinct numbers  $a_1, a_2, \dots, a_\nu \in \{1, 2, \dots, N - f + \nu - 1\}$ , where  $0 < a_1 < a_2 < \dots < a_\nu \leq N - f + \nu - 1$ , such that for every  $i \in \{1, 2, \dots, \nu\}$ ,*

- (i)  $P_i^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i-1} - 1, a_i, a_{i+1}, \dots, a_\nu)$  is  $(\sigma(i), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\})$ -valent;
- (ii)  $P_i^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i-1} - 1, a_i, a_{i+1}, \dots, a_\nu)$  is not  $(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\})$ -valent, for any  $1 \leq j < i$ ;
- (iii) if  $P_i^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i-1} - 1, a_i, a_{i+1}, \dots, a_\nu)$  is  $(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i-1)}, C_{\sigma(j)}\})$ -valent for some  $i < j \leq \nu$ , then  $v_{\sigma(i)} \prec v_{\sigma(j)}$ .

In our proof of Lemma 6.10, we use Lemmas 6.11, 6.12, 6.13 and 6.14 which are stated and proved below. The next two lemmas state properties of the point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ .

**Lemma 6.11.** *If  $a_1 = N - f$ , then, for every permutation  $\sigma$  and integers  $a_2, a_3, \dots, a_\nu \in \{N - f, N - f + 1, \dots, N - f + \nu - 1\}$  where  $a_1 \leq a_2 \leq \dots \leq a_\nu$ , the point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(i, \mathcal{C}_w - \{C_i\})$ -valent for some  $i \in \{1, 2, \dots, \nu\}$ .*

*Proof.* Let  $a_1 = N - f$ . Consider an execution  $\beta$  that extends  $\alpha_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  as follows. Note that at point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , all the value-dependent messages from the clients have been delivered to the first  $N - f$  servers. Therefore, for every client  $C_i, i \in \{1, 2, \dots, \nu\}$ , its write operation  $\pi_i$  has sent its value-dependent messages, and the channels from the client to the servers have delivered all the value-dependent messages in the execution; all the send actions enabled on behalf of operation  $\pi_i$  are value-independent actions. In the remainder of the execution  $\beta$ , the last  $f$  servers do not take any steps. The clients  $C_1, C_2, \dots, C_\nu$ , their channels and the first  $N - f$  servers perform their actions in a fair manner. Since algorithm  $A$  ensures that every write operation terminates in a fair execution so long as the number of server failures is no bigger than  $f$  and the number of active write clients is no bigger than  $\nu$ , we know that a write operation  $\pi_j$  completes in  $\beta$  for some  $j \in \{1, 2, \dots, \nu\}$ . After the completion of a write operation  $\pi_j$ , the write clients and the channels from the write clients stop performing actions. A read operation  $\pi_r$  begins at a reader. The reader and the first  $N - f$  servers perform actions in a fair manner until read operation  $\pi_r$  terminates. Because read operation  $\pi_r$  begins after the termination of  $\pi_j$ , and because the algorithm satisfies weak regularity, the operation returns  $v_i$  for some  $i \in \{1, 2, \dots, \nu\}$ . The execution  $\beta$  finishes after the termination of read  $\pi_r$ . Thus we have created an execution  $\beta$ , which is an extension of  $\alpha_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  such that all the clients and their channels only take value-independent output actions, and a read operation returns  $v_i$ . Therefore the point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $i$ -valent, and thus is also  $(i, \mathcal{C}_w - \{C_i\})$ -valent.  $\square$

**Lemma 6.12.** *If point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(j, \mathcal{C}_w - \{C_j\})$ -valent for some permutation  $\sigma$  and positive integers  $j, a_1, a_2, \dots, a_\nu$  where  $0 \leq a_1 \leq a_2 \leq \dots \leq a_\nu \leq N - f + \nu - 1$ , then  $a_1 \geq 1$ .*

*Proof.* To show that  $a_1 \geq 1$ , assume to the contrary that  $a_1 = 0$ . Because the point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(j, \mathcal{C}_w - \{C_j\})$ -valent, there is an execution  $\beta$  which extends  $\alpha_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  such that, client  $C_j$  stops acting at point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , and a read operation  $\pi_r$  begins and returns  $v_j$ . Now, because  $a_1 = 0$ , we note that point  $P_1^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is the same as point  $P_0^{\vec{v}}$ .

Consider the value vector  $\vec{u} = (u_0, u_1, u_2, \dots, u_\nu)$ , where for  $i \in \{0, 1, \dots, \nu\} - \{j\}$ , we have  $u_i = v_i$ , and  $v_j \notin \{u_0, u_1, \dots, u_\nu\}$ . Note that at point  $P_0^{\vec{v}}$ , the state of every server, channel from a server to a server or a client, and client in  $\mathcal{C}_w - \{C_j\}$  is the same as its state at point  $P_0^{\vec{u}}$ . Consider an execution  $\beta'$  which extends  $\alpha_0^{\vec{u}}$  as follows. Starting from point  $P_0^{\vec{u}}$ , every component takes the same steps as the component takes starting from point  $P_0^{\vec{v}}$  in  $\beta$ . Note that client  $C_j$  takes only value-independent actions after  $P_0^{\vec{v}}$  in  $\beta$ . Because Lemma 6.9 implies that the state of every component except client  $C_j$ , and the metadata component of  $C_j$  is the at point  $P_0^{\vec{u}}$  in  $\beta'$  as at point  $P_0^{\vec{v}}$  in  $\beta$ , execution  $\beta'$  is an execution of the algorithm. In execution  $\beta'$ , read operation  $\pi_r$  returns  $v_j$  which does not belong to  $\{u_0, u_1, \dots, u_\nu\}$ . Therefore execution  $\beta'$  violates weak regularity. Therefore  $a_1 \geq 1$ .  $\square$

The next lemma shows that, informally, the valency of the point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  does not depend on value  $a_{i+1}, \dots, a_\nu, \sigma(i+1), \dots, \sigma(\nu)$ .

**Lemma 6.13.** *Suppose that a point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\})$ -valent for some permutation  $\sigma$  and integers  $j, a_1, \dots, a_\nu$  such that  $0 \leq a_1 \leq a_2 \leq \dots \leq a_\nu \leq N - f + \nu - 1$ . Then the point  $P_i^{\vec{v}}(\bar{\sigma}, a_1, a_2, \dots, a_i, \bar{a}_{i+1}, \bar{a}_{i+2}, \dots, \bar{a}_\nu)$  is  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\})$ -valent for every set of integers  $\bar{a}_{i+1}, \bar{a}_{i+2}, \dots, \bar{a}_\nu$  where  $a_i \leq \bar{a}_{i+1} \leq \bar{a}_{i+2} \leq \dots \leq \bar{a}_\nu \leq N - f + \nu - 1$ , and every  $\bar{\sigma}$  where  $\bar{\sigma}(l) = \sigma(l), 1 \leq l \leq i$ .*

*Proof of Lemma 6.13.* Let  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  be  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\})$ -valent. Therefore, there exists an execution  $\beta$  that extends  $\alpha_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  such that, after point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  the clients and the channels from the clients in  $C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}$  do not take any value-dependent actions, and there is a read operation  $\pi_r$  that begins and returns value  $v_j$ . Now we construct execution  $\beta'$  which extends

$$\alpha_i^{\vec{v}}(\bar{\sigma}, a_1, a_2, \dots, a_i, \bar{a}_{i+1}, \bar{a}_{i+2}, \dots, \bar{a}_\nu)$$

as follows. Note that at point  $P_i^{\vec{v}}(\bar{\sigma}, a_1, a_2, \dots, a_i, \bar{a}_{i+1}, \bar{a}_{i+2}, \dots, \bar{a}_\nu)$  every server, channel and client is at the same state as it is at point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  in  $\alpha_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ . In execution  $\beta'$ , every component performs the same actions as the component does in  $\beta$  starting from point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ . Therefore, after point  $P_i^{\vec{v}}(\bar{\sigma}, a_1, a_2, \dots, a_i, \bar{a}_{i+1}, \bar{a}_{i+2}, \dots, \bar{a}_\nu)$  the clients and the channels from the clients in  $C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}$  do not take any value-dependent actions, and there is a read operation  $\pi_r$  that begins and returns value  $v_j$ . Therefore the point

$$P_i^{\vec{v}}(\bar{\sigma}, a_1, a_2, \dots, a_i, \bar{a}_{i+1}, \bar{a}_{i+2}, \dots, \bar{a}_\nu)$$

is  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i)}\})$ -valent.  $\square$

The next lemma, informally, shows that if a point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  has sufficient information of a value  $v_j$ , then an earlier point  $P_k^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ ,  $k \leq i$ , already has sufficient information of  $v_j$  if we allow some extra clients to take value-dependent actions.

**Lemma 6.14.** *Let  $1 \leq k \leq l \leq i \leq \nu$ . If a point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(l)}\})$ -valent for some permutation  $\sigma$  and integers  $j, a_1, \dots, a_\nu$  such that  $0 \leq a_1 \leq a_2 \leq \dots \leq a_\nu \leq N - f + \nu - 1$ , then for every  $k \leq i$ , the point  $P_k^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(k)}\})$ -valent.*

*Proof.* Let  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  be  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(l)}\})$ -valent. Therefore, there exists an execution  $\beta$  that extends  $\alpha_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  such that, after point  $P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  the clients and the channels from the clients in  $C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(l)}$  do not take any value-dependent actions, and there is a read operation  $\pi_r$  that begins and returns value  $v_j$ . Note that for  $k \leq i$ , execution  $\beta$  is an extension of  $\alpha_k^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , where, after point  $P_k^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$ , the clients in  $\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(k)}\}$  and the channels from these clients do not perform value-dependent actions, read operation  $\pi_r$  that begins and returns value  $v_j$ . Therefore the point  $P_k^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  is  $(j, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(k)}\})$ -valent.  $\square$

We are now ready to prove our key lemma, Lemma 6.10.

*Proof of Lemma 6.10.* We begin by choosing  $a_1, \sigma(1)$ . From Lemma 6.11, we know that the set

$$\mathcal{A}_1 = \left\{ (\bar{a}_1, i) : \begin{array}{l} \text{there exists an integer } i \in \{1, 2, \dots, \nu\}, \text{ a permutation } \bar{\sigma}, \text{ and} \\ \text{integers } \bar{a}_1, \dots, \bar{a}_\nu \text{ where } 0 \leq \bar{a}_1 \leq \bar{a}_2 \leq \dots \leq \bar{a}_\nu \leq N - f + \nu - 1 \\ \text{such that } P_1^{\bar{\sigma}}(\bar{\sigma}, \bar{a}_1, \bar{a}_2, \dots, \bar{a}_\nu) \\ \text{is } (i, \mathcal{C}_w - \{C_i\})\text{-valent} \end{array} \right\} \quad (1)$$

is non-empty. In particular, Lemma 6.11 implies that tuple  $(N - f, i)$  belongs to  $\mathcal{A}_1$  for some integer  $i \in \{1, 2, \dots, \nu\}$ . We choose  $a_1$  to be the smallest integer such that  $(a_1, i)$  belongs to the set  $\mathcal{A}_1$  for some  $i$ <sup>7</sup>, that is

$$a_1 = \min\{a : \text{there exists } j \text{ such that } (a, j) \in \mathcal{A}_1\}.$$

Note that Lemma 6.12 implies that  $a_1 \geq 1$ . We let

$$\sigma(1) = \arg \min_{\{j : (a_1, j) \in \mathcal{A}_1\}} v_j,$$

where the minimum is according to the ordering  $\prec$ .

Next, we provide a procedure that recursively chooses  $a_{i_0+1}$  and  $\sigma(i_0 + 1)$  given  $a_1, a_2, \dots, a_{i_0}$ , and  $\sigma(1), \sigma(2), \dots, \sigma(i_0)$ , for any  $i_0 \in \{1, 2, \dots, \nu - 1\}$ . Our choice of  $a_{i_0+1}$  will satisfy  $a_{i_0} < a_{i_0+1} \leq N - f + i_0 - 1$ .

Let

$$\mathcal{A}_{i_0+1} = \left\{ (\bar{a}_{i_0+1}, i) : \begin{array}{l} \text{there exists a permutation } \bar{\sigma}, \text{ where } \bar{\sigma}(j) = \sigma(j) \text{ for } j \leq i_0, \text{ and} \\ \text{integers } \bar{a}_{i_0+1}, \bar{a}_{i_0+2}, \dots, \bar{a}_\nu \text{ where } a_{i_0} \leq \bar{a}_{i_0+1} \leq \bar{a}_{i_0+2} \leq \dots \leq \bar{a}_\nu \leq N - f + \nu - 1 \\ \text{and an integer } i \in \{1, 2, \dots, \nu\} - \{\sigma(1), \sigma(2), \dots, \sigma(i_0)\} \\ \text{such that } P_{i_0+1}^{\bar{\sigma}}(\bar{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, \bar{a}_{i_0+1}, \bar{a}_{i_0+2}, \dots, \bar{a}_\nu) \\ \text{is } (i, \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}, C_i\})\text{-valent} \end{array} \right\} \quad (2)$$

We show that  $\mathcal{A}_{i_0+1}$  is non-empty. We choose  $a_{i_0+1}$  to be the smallest integer such that  $(a_{i_0+1}, i)$  belongs to the set  $\mathcal{A}_{i_0+1}$  for some  $i$ . We let

$$\sigma(i_0 + 1) = \arg \min_{\{j : (a_{i_0+1}, j) \in \mathcal{A}_{i_0+1}\}} v_j,$$

where the minimum above is taken according to the ordering  $\prec$ .

To complete the proof of Lemma 6.10, we show that

- (a)  $\mathcal{A}_{i_0+1}$  is non-empty, and  $a_{i_0+1} > a_{i_0}$ ,  $1 \leq i_0 \leq \nu - 1$ ;
- (b)  $P_{i_0+1}^{\bar{\sigma}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is  $(\sigma(i_0 + 1), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0+1)}\})$ -valent,  $0 \leq i_0 \leq \nu - 1$ ;
- (c)  $P_{i_0+1}^{\bar{\sigma}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is not  $(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0+1)}\})$ -valent, for any  $j < i_0 + 1$ ,  $1 \leq i_0 \leq \nu - 1$ ;
- (d) if  $P_{i_0+1}^{\bar{\sigma}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is

$$(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, \dots, C_{\sigma(i_0)}, C_{\sigma(j)}\})\text{-valent}$$

for some  $j > i_0 + 1$ , then  $v_{\sigma(i_0+1)} \prec v_{\sigma(j)}$ ,  $0 \leq i_0 \leq \nu - 2$ .

*Proof of (a):*

We show (a) by showing that there is some integer  $i$  such that  $(N - f + i_0, i)$  belongs to  $\mathcal{A}_{i_0+1}$ . More specifically, we will show that for any arbitrary permutation  $\bar{\sigma}$  which satisfies  $\bar{\sigma}(j) = \sigma(j)$  for  $1 \leq j \leq i_0$ , the point  $P_{i_0+1}^{\bar{\sigma}}(\bar{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, N - f + i_0, N - f + i_0 + 1, \dots, N - f + \nu - 1)$  is  $(i, \mathcal{C}_w -$

---

<sup>7</sup>Informally,  $a_1$  may be viewed as the smallest number such that the first  $a_1$  servers contains ‘‘sufficient information’’ of some value  $v_i$ , given the information of all other values.

$\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}, C_i\}$ -valent for some  $i \in \{1, 2, \dots, \nu\} - \{\sigma(1), \sigma(2), \dots, \sigma(i_0)\}$ . To show this, we construct an execution  $\beta$ , which is an extension of  $\alpha_{i_0+1}^{\vec{v}}(\vec{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, N - f + i_0, N - f + i_0 + 1, \dots, N - f + \nu - 1)$  as follows. After point  $P_{i_0+1}^{\vec{v}}(\vec{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, N - f + i_0, N - f + i_0 + 1, \dots, N - f + \nu - 1)$  in  $\beta$ , the write clients in  $\{C_{\vec{\sigma}(1)}, C_{\vec{\sigma}(2)}, \dots, C_{\vec{\sigma}(i_0)}\}$ , and the channels from these write clients to the non-failed servers do not send or deliver value-dependent messages. The clients in  $\mathcal{C}_w - \{C_{\vec{\sigma}(1)}, C_{\vec{\sigma}(2)}, \dots, C_{\vec{\sigma}(i_0)}\}$ , the channels from these clients, the non-failed servers, and the channels between the servers continue taking actions in a fair manner. Note that algorithm A guarantees that in a fair execution where the number of server failures is no bigger than  $f$  and the number of active write clients is no bigger than  $\nu$ , every write operation terminates. From the perspective of clients in  $\mathcal{C}_w - \{C_{\vec{\sigma}(1)}, C_{\vec{\sigma}(2)}, \dots, C_{\vec{\sigma}(i_0)}\}$ , the execution  $\beta$  is indistinguishable from a fair execution. Therefore some operation in  $\{\pi_{\vec{\sigma}(i_0+1)}, \pi_{\vec{\sigma}(i_0+2)}, \dots, \pi_{\vec{\sigma}(\nu)}\}$  terminates in  $\beta$ . After the termination of an operation  $\pi_j, j \in \{\vec{\sigma}(i_0 + 1), \vec{\sigma}(i_0 + 2), \dots, \vec{\sigma}(\nu)\}$ , all the write operations and their channels stop taking actions. A read operation  $\pi_r$  begins at a read client. The reader, the channels from and to the reader, and the non-failed servers perform actions in a fair manner until the read operation  $\pi_r$  terminates. After the termination of  $\pi_r$  the execution  $\beta$  ends. Because the algorithm satisfies regularity and because write operation  $\pi_j$  has terminated, the read operation  $\pi_r$  returns value  $v_i$  for some  $i \in \{1, 2, \dots, \nu\}$ .

We show that, in fact,  $\pi_r$  returns  $v_i$  for some  $i \in \{1, 2, \dots, \nu\} - \{\sigma(1), \sigma(2), \dots, \sigma(i_0)\}$ . Assume the contrary, that is, assume that the read operation  $\pi_r$  returns  $v_{\sigma(k)}$  for  $k \in \{1, 2, \dots, i_0\}$ . The existence of execution  $\beta$  implies that the point  $P_{i_0+1}^{\vec{v}}(\vec{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, N - f + i_0, N - f + i_0 + 1, \dots, N - f + \nu - 1)$  is  $(\sigma(k), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent. Lemma 6.14 implies that  $P_k^{\vec{v}}(\vec{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, N - f + i_0, N - f + i_0 + 1, \dots, N - f + \nu - 1)$  is  $(\sigma(k), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(k)}\})$ -valent. Therefore  $(a_k - 1, \sigma(k)) \in \mathcal{A}_k$ . This however contradicts the fact that we choose  $a_k$  to be the smallest element such that  $(a_k, j)$  is  $\mathcal{A}_k$  for some  $j \in \{1, 2, \dots, \nu\} - \{\sigma(1), \sigma(2), \dots, \sigma(k - 1)\}$ . Therefore, it cannot be that  $\pi_r$  returns  $v_{\sigma(k)}$  for some  $k \in \{1, 2, \dots, i_0\}$ . Therefore,  $\mathcal{A}_{i_0+1}$  is non-empty.

We now show that  $a_{i_0+1} > a_{i_0}$ . We know  $a_{i_0+1} \geq a_{i_0}$ . To show that  $a_{i_0+1} > a_{i_0}$ , assume to the contrary that  $a_{i_0+1} = a_{i_0}$ . Because  $a_{i_0+1} \in \mathcal{A}_{i_0+1}$  and because we assume  $a_{i_0+1} = a_{i_0}$ , point  $P_{i_0+1}^{\vec{v}}(\vec{\sigma}, a_1 - 1, \dots, a_{i_0} - 1, a_{i_0}, \bar{a}_{i_0+1}, \dots, \bar{a}_\nu)$  is  $(\sigma(i_0 + 1), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}, C_{\sigma(i_0+1)}\})$ -valent. By Lemma 6.14 this implies that point  $P_{i_0}^{\vec{v}}(\vec{\sigma}, a_1 - 1, \dots, a_{i_0} - 1, a_{i_0}, \bar{a}_{i_0+1}, \dots, \bar{a}_\nu)$  is  $(\sigma(i_0 + 1), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent. Therefore,  $(a_{i_0} - 1, \sigma(i_0 + 1)) \in \mathcal{A}_{i_0}$ , and contradicts the fact that we choose  $a_{i_0}$  to be the smallest element such that  $(a_{i_0}, j)$  is in  $\mathcal{A}_{i_0}$  for some  $j \in \{1, 2, \dots, \nu\} - \{\sigma(1), \sigma(2), \dots, \sigma(i_0 - 1)\}$ .

*Proof of (b):*

Because  $(a_{i_0}, \sigma(i_0)) \in \mathcal{A}_{i_0}, 1 \leq i_0 \leq \nu$  there are distinct integers  $\bar{a}_{i_0+1}, \bar{a}_{i_0+2}, \dots, \bar{a}_\nu$ , and a permutation  $\vec{\sigma}$  such that

- $\vec{\sigma}(j) = \sigma(j)$  for  $1 \leq j \leq i_0$
- $a_{i_0} \leq \bar{a}_{i_0+1} \leq \bar{a}_{i_0+2} \leq \dots \bar{a}_\nu \leq N - f + \nu - 1$

such that the point

$$P_{i_0}^{\vec{v}}(\vec{\sigma}, a_1 - 1, a_2 - 1, \dots, a_{i_0-1} - 1, a_{i_0}, \bar{a}_{i_0+1}, \bar{a}_{i_0+2}, \dots, \bar{a}_\nu)$$

is  $(\sigma(i_0), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent. Using Lemma 6.13, we conclude that the point

$$P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0-1} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$$

is  $(\sigma(i_0), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent.

*Proof of (c):*

If point  $P_{i_0+1}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is  $(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0+1)}\})$ -valent for  $j < i_0 + 1$ , then by Lemma 6.14, the point  $P_j^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is  $(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(j)}\})$ -valent. However, this implies that  $(a_j - 1, \sigma(j)) \in \mathcal{A}_j$ , which contradicts the fact that  $a_j$  is the smallest integer among all the integers such that  $(a_j, k) \in \mathcal{A}_j$ . Therefore  $P_{i_0+1}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is not  $(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent, for any  $j < i_0$ .

*Proof of (d):* If  $P_{i_0+1}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0+1}, a_{i_0+2}, \dots, a_\nu)$  is

$$(\sigma(j), \mathcal{C}_w - \{C_{\sigma(1)}, \dots, C_{\sigma(i_0)}, C_{\sigma(j)}\})\text{-valent}$$

for some  $j > i_0 + 1$ , then  $(a_{i_0+1}, \sigma(j))$  belongs to  $\mathcal{A}_{i_0+1}$ . Because  $(a_{i_0+1}, \sigma(i_0 + 1))$  also belongs to  $\mathcal{A}_{i_0+1}$  and because we chose

$$\sigma(i_0 + 1) = \arg \min_{\{(a_{i_0+1}, k) \in \mathcal{A}_{i_0+1}\}} v_k,$$

we have  $v_{\sigma(i_0+1)} \prec v_{\sigma(j)}$ . □

#### 6.4.4 Proof of Theorem 6.5

Recall that  $\mathcal{S}_n$  represents the set of possible server states of the  $n$ th server. Let  $\vec{S}_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$  denote the  $N - f + \nu - 1$  dimensional vector in the set  $\prod_{n=1}^{N-f+\nu-1} \mathcal{S}_n$ , whose  $j$ th component denotes the state of server  $j$  at point

$$P_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu)$$

in execution

$$\alpha_i^{\vec{v}}(\sigma, a_1, a_2, \dots, a_\nu).$$

For a given vector  $\vec{v} = (v_0, v_1, \dots, v_\nu) \in \mathcal{V}^\nu$ , let permutation  $\sigma^{\vec{v}}$  and distinct integers  $a_1^{\vec{v}}, \dots, a_\nu^{\vec{v}}$  satisfy the conditions of Lemma 6.10 as per a total order  $\prec$  on the set  $\mathcal{V}$ .

Let  $v_0 \in \mathcal{V}$  be the initial value, and

$$\mathcal{V}_0 = \{(v_0, v_1, \dots, v_\nu) : v_1, v_2, \dots, v_\nu \in \mathcal{V} - \{v_0\} \text{ are distinct}\}.$$

We show that there is a one-to-one mapping from tuples of the form

$$(\sigma^{\vec{v}}, a_1^{\vec{v}}, a_2^{\vec{v}}, \dots, a_\nu^{\vec{v}}, \vec{S}_\nu^{\vec{v}}(\sigma^{\vec{v}}, a_1^{\vec{v}}, a_2^{\vec{v}}, \dots, a_\nu^{\vec{v}}))$$

to the vectors in  $\mathcal{V}_0$ . Specifically, if  $\vec{u}, \vec{v} \in \mathcal{V}_0$  and  $\vec{u} \neq \vec{v}$ , we show that

$$\begin{aligned} &(\sigma^{\vec{u}}, a_1^{\vec{u}}, a_2^{\vec{u}}, \dots, a_\nu^{\vec{u}}, \vec{S}_\nu^{\vec{u}}(\sigma^{\vec{u}}, a_1^{\vec{u}}, a_2^{\vec{u}}, \dots, a_\nu^{\vec{u}})) \\ &\neq (\sigma^{\vec{v}}, a_1^{\vec{v}}, a_2^{\vec{v}}, \dots, a_\nu^{\vec{v}}, \vec{S}_\nu^{\vec{v}}(\sigma^{\vec{v}}, a_1^{\vec{v}}, a_2^{\vec{v}}, \dots, a_\nu^{\vec{v}})). \end{aligned} \quad (3)$$

The one-to-one mapping implies that the cardinality of  $\mathcal{P} \times \{1, 2, \dots, N - f + \nu - 1\}^\nu \times \prod_{n=1}^{N-f+\nu-1} \mathcal{S}_n$  must be no smaller than  $|\mathcal{V}_0|$ , where  $\mathcal{P}$  represents the set of all permutations on  $\{1, 2, \dots, \nu\}$ . This implies that

$$(\nu!) \cdot (N - f + \nu - 1)^\nu \cdot \prod_{n=1}^{N-f+\nu-1} |\mathcal{S}_n| \geq |\mathcal{V}_0| = \binom{|\mathcal{V}| - 1}{\nu}$$

which implies the statement of the theorem.

To complete the theorem, we show the relation stated in (3). We provide a proof by contradiction. Consider two distinct vectors  $\vec{u} = (v_0, u_1, \dots, u_\nu)$  and  $\vec{v} = (v_0, v_1, \dots, v_\nu)$  which violate (3). Let  $\sigma = \sigma^{\vec{u}} = \sigma^{\vec{v}}$  and  $a_i = a_i^{\vec{u}} = a_i^{\vec{v}}$  for  $i \in \{1, 2, \dots, \nu\}$ .

Since  $\vec{u} \neq \vec{v}$ , we know that there exists an index  $i \in \{1, 2, \dots, \nu\}$  such that  $u_i \neq v_i$ . Let  $i_0$  be the largest element of  $\{j : u_{\sigma(j)} \neq v_{\sigma(j)}\}$ . Note that if  $j > i_0$ , we have  $u_{\sigma(j)} = v_{\sigma(j)}$ . Also,  $u_{\sigma(i_0)} \neq v_{\sigma(i_0)}$ . This implies that either  $u_{\sigma(i_0)} \prec v_{\sigma(i_0)}$  or  $v_{\sigma(i_0)} \prec u_{\sigma(i_0)}$ . Without loss of generality, we assume that  $v_{\sigma(i_0)} \prec u_{\sigma(i_0)}$ . We next use Lemma 6.10 to show that  $u_{\sigma(i_0)} = v_{\sigma(i_0)}$  which is a contradiction.

Because the point  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  is  $(\sigma(i_0), \mathcal{C}_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent, there exists an execution  $\beta'$  which extends

$$\alpha_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0-1} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu),$$

such that after  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  the clients in  $\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\}$  and the channels from these clients do not take value-dependent actions and there is a read operation that begins and returns  $v_{\sigma(i_0)}$ .

We compare the component states at two points: point  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  in execution  $\beta'$  and point  $P_{i_0}^{\vec{u}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  in execution  $\alpha_{i_0}^{\vec{u}}(\sigma, a_1 - 1, a_2 -$

$1, \dots, a_{i_0-1} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu$ ). Because  $\vec{S}_\nu^{\vec{u}} = \vec{S}_\nu^{\vec{v}}$ , the state of a server at the first point is the same as its state at the second point. From Lemma 6.9, the metadata components of the state of any write client and the state of every channel from a server to a server or client are the same at both points. Finally, because  $u_{\sigma(j)} = v_{\sigma(j)}$  for all  $j > i_0$ , and the clients in  $C_w - \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\}$  have the same state at both points.

We now create an execution  $\beta$  that extends  $\alpha_{i_0}^{\vec{u}}(\sigma, a_1, a_2, \dots, a_\nu)$  as follows. Starting at point  $P_{i_0}^{\vec{u}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$ , every client, server, and channel takes the same steps in  $\beta$  as it takes starting from point  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  in  $\beta'$ . Note that every component except the clients in  $\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\}$  and the channels from these clients have the same state at point  $P_{i_0}^{\vec{u}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  in  $\beta$  as at point  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  in  $\beta'$ . The metadata components of the states of the clients in  $\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\}$  and the metadata messages in the channels from these clients are the same at the two points. Because clients in  $\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\}$  and the channels from these clients only take value-independent output actions after point  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$  in  $\beta'$ ,  $\beta$  is a valid execution of the algorithm  $A$ . Thus, we have created an execution  $\beta$  which is an extension of  $\alpha_{i_0}^{\vec{u}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$ , where, after point  $P_{i_0}^{\vec{v}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$ , the clients in  $\{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\}$  and the channels from these clients do not take any value-dependent actions, and a read operation begins and returns  $v_{\sigma(i_0)}$ . Because the algorithm is weakly regular, we must have  $v_{\sigma(i_0)} \in \{v_0, u_1, u_2, \dots, u_\nu\}$ . Furthermore,  $v_{\sigma(i_0)} \neq v_0$  since we assume that the components of  $\vec{v}$  are distinct. So, there exists an integer  $j_0$  in  $\{1, 2, \dots, \nu\}$  such that  $v_{\sigma(i_0)} = u_{\sigma(j_0)}$ .

The existence of execution  $\beta$  implies that the point  $P_{i_0}^{\vec{u}}(\sigma, a_1 - 1, a_2 - 1, \dots, a_{i_0} - 1, a_{i_0}, a_{i_0+1}, \dots, a_\nu)$ , is  $(\sigma(j_0), \{C_{\sigma(1)}, C_{\sigma(2)}, \dots, C_{\sigma(i_0)}\})$ -valent. Statement (ii) of Lemma 6.10 implies that  $j_0 \geq i_0$ . Statement (iii) of Lemma 6.10 implies that if  $j_0 > i_0$ , then  $u_{\sigma(i_0)} \prec u_{\sigma(j_0)}$ . However, we started with the assumption that  $u_{\sigma(j_0)} = v_{\sigma(i_0)} \prec u_{\sigma(i_0)}$ . Therefore, it cannot happen that  $j_0 > i_0$ , and we conclude that  $j_0 = i_0$ ,  $v_{\sigma(i_0)} = u_{\sigma(i_0)}$ . This contradicts our assumption that  $v_{\sigma(i_0)} \neq u_{\sigma(i_0)}$ . Therefore, (3) must be true. This completes the proof.

## 6.5 Conjecture related to Theorem 6.5

The assumptions of Theorem 6.5 do not apply to some algorithms [2, 15]. These algorithms send value-dependent messages in two phases. In one of the two phases, the algorithms send erasure coded elements corresponding to the value. In the other phase where value-dependent messages are sent, a hash of the value is sent. The hashes are used for verification of the client's integrity, which is important in [2, 15] as the algorithms in these references handle Byzantine failures. We believe that it may be possible to generalize our result of Theorem 6.5 with Assumption 3 (b) modified as follows:

- the algorithm has a bounded number of phases, and
- there is at most one phase where a value-dependent message of size  $\Theta(|\mathcal{V}|)$  is sent.

The above restrictions imply that, even if there is more than one phase where value dependent messages are sent, the value-dependent messages in the additional phases do not carry much information about the value. The above restrictions would cover algorithms of [2, 15], and we conjecture that the lower bound of Corollary 6.6 bound still applies.

## 7 Concluding Remarks

This paper was motivated by the following open question (see Section 2): Does there exist an atomic shared memory emulation algorithm whose storage cost is smaller than  $\nu \frac{N}{N-f} \log_2 |\mathcal{V}|$ , where  $\nu$  represents the number of active write operations? This question remains open. The insight obtained by our bounds in conjunction with the result of [23] is summarized here. If there is an algorithm whose storage cost is  $g(\nu, N, f) \log_2 |\mathcal{V}| + o(\log_2 |\mathcal{V}|)$ , where  $g(\nu, N, f)$  is some real-valued function of parameters  $\nu, N, f$  then

- $g(\nu, N, f) \geq \frac{2N}{N-f+2}$ ;

- If  $g(\nu, N, f) < \frac{\nu N}{N-f+\nu-1}$ , then
  - the writer sends its value in multiple phases to the servers, or
  - the writer's state may not separate the value and the metadata, or
  - during a write operation, the writer can take non-black box actions;
- If, for a given values of parameters  $N, f$ , we have  $g(\nu, N, f) < f + 1$  for all values of  $\nu$ , then, in certain executions, the servers store symbols which jointly encode values across different versions.



## References

- [1] M. K. Aguilera, R. Janakiraman, and L. Xu. Using erasure codes efficiently for storage in a distributed system. In *Dependable Systems and Networks, 2005. DSN 2005. Proceedings. International Conference on*, pages 336–345. IEEE, 2005.
- [2] E. Androulaki, C. Cachin, D. Dobre, and M. Vukolić. Erasure-coded Byzantine storage with separate metadata. In *Principles of Distributed Systems*, pages 76–90. Springer, 2014.
- [3] H. Attiya, A. Bar-Noy, and D. Dolev. Sharing memory robustly in message-passing systems. In *Proceedings of the ninth annual ACM symposium on principles of distributed computing*, PODC '90, pages 363–375. ACM, 1990.
- [4] C. Cachin and S. Tessaro. Optimal resilience for erasure-coded Byzantine distributed storage. In *Dependable Systems and Networks, 2006. DSN 2006. International Conference on*, pages 115–124. IEEE, 2006.
- [5] V. Cadambe, N. Lynch, M. Medard, and P. Musial. A coded shared atomic memory algorithm for message passing architectures. In *2014 IEEE 13th International Symposium on Network Computing and Applications (NCA)*, pages 253–260, 2014.
- [6] V. R. Cadambe, N. Lynch, M. Medard, and P. Musial. A coded shared atomic memory algorithm for message passing architectures. 2014. MIT CSAIL Technical Report MIT-CSAIL-TR-2014-015, <http://hdl.handle.net/1721.1/88551>; Also available on <http://arxiv.org/abs/1407.4167>.
- [7] Y. Cassuto. What can coding theory do for storage systems? *ACM SIGACT News*, 44(1):80–88, 2013.
- [8] G. Chockler, D. Dobre, A. Shraer, and A. Spiegelman. Space bounds for reliable multi-writer data store: Inherent cost of read/write primitives. 2015. arXiv pre-print, available on <http://arxiv.org/abs/1508.03762>.
- [9] G. Chockler, R. Guerraoui, and I. Keidar. Amnesic distributed storage. In *Distributed Computing*, pages 139–151. Springer, 2007.
- [10] A. G. Dimakis, P. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran. Network coding for distributed storage systems. *IEEE Transactions on Information Theory*, 56(9):4539–4551, 2010.
- [11] D. Dobre, G. Karame, W. Li, M. Majuntke, N. Suri, and M. Vukolić. PoWerStore: proofs of writing for efficient and robust storage. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 285–298. ACM, 2013.
- [12] P. Dutta, R. Guerraoui, and R. R. Levy. Optimistic erasure-coded distributed storage. In *Distributed Computing*, pages 182–196. Springer, 2008.
- [13] R. Fan and N. Lynch. Efficient replication of large data objects. In *In Proceedings of the 17th International Symposium on Distributed Computing (DISC)*, pages 75–91, 2003.
- [14] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382, 1985.
- [15] J. Hendricks, G. R. Ganger, and M. K. Reiter. Low-overhead Byzantine fault-tolerant storage. *ACM SIGOPS Operating Systems Review*, 41(6):73–86, 2007.
- [16] M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12:463–492, July 1990.
- [17] L. Lamport. On interprocess communication. Part I: basic formalism. *Distributed Computing*, 2(1):77–85, 1986.
- [18] S. Lin and D. J. Costello. *Error control coding, Second edition*. Prentice-Hall, Inc., 2004.

- [19] D. A. Patterson, G. Gibson, and R. H. Katz. *A case for redundant arrays of inexpensive disks (RAID)*, volume 17. ACM, 1988.
- [20] R. Roth. *Introduction to coding theory*. Cambridge University Press, 2006.
- [21] Y. Saito, S. Frølund, A. Veitch, A. Merchant, and S. Spence. FAB: building distributed enterprise disk arrays from commodity components. *ACM SIGOPS Operating Systems Review*, 38(5):48–58, 2004.
- [22] C. Shao, J. L. Welch, E. Pierce, and H. Lee. Multiwriter consistency conditions for shared memory registers. *SIAM Journal on Computing*, 40(1):28–62, 2011.
- [23] A. Spiegelman, Y. Cassuto, G. Chockler, and I. Keidar. Space bounds for reliable storage: Fundamental limits of coding. In *Proceedings of the 2016 ACM symposium on principles of distributed computing*, PODC '16. ACM, 2016.
- [24] Z. Wang and V. R. Cadambe. Multi-version coding - an information theoretic perspective of consistent distributed storage. arxiv pre-print, available on <http://arxiv.org/abs/1506.00684>.
- [25] Z. Wang and V. R. Cadambe. Multi-version coding in distributed storage. In *2014 IEEE International Symposium on Information Theory (ISIT)*, 2014.

## A Discussion on the Storage Scheme Assumption of [23]

For the sake of technical clarity and completeness, we provide a discussion on the storage scheme assumption of [23]. Reference [23] assumes that every stored bit is associated uniquely with a value of a write operation. However, this assumption is restrictive, and the storage cost lower bound proof of [23] may not be applicable for arbitrary storage schemes.

In fact, a critical idea in the proof of [23] is the following. If at a point  $P$  in an execution of an algorithm  $A$ ,

- no value from  $\{v_1, v_2, \dots, v_m\} \subset \mathcal{V}$  is returnable from a set of servers,
- if there is a value  $v_i \in \{v_1, v_2, \dots, v_m\}$  such that no server stores a single bit of the value, and
- after a single step of the execution, some value which is not necessarily  $\{v_1, v_2, \dots, v_m\} \subset \mathcal{V}$  is returnable,

then the number of bits stored in a server must increase by  $\log_2 |\mathcal{V}|$  bits in the step. However, this is true only for specific storage schemes, but may not be true for arbitrary storage schemes. We show this via a counter-example.

Let  $\mathcal{V}$  be a finite field of  $2^m$  elements for some integer  $m$ . Note that every element of  $\mathcal{V}$  is an  $m$ -bit vector over the binary base field. Let  $v_1, v_2, v_3 \in \mathcal{V}$  be three versions of the data object associated with three write operations in an execution of an algorithm. Suppose that, in some algorithm  $A$ , because of structure of the server protocol, there are two servers which both store  $v_1 + v_2 + v_3$  at some point  $P$  of an execution. It is impossible to associate a bit stored at the servers with a single value. Note that a reader which accesses only the two servers cannot recover even a single bit of  $v_1, v_2, v_3$ .

For argument's sake, suppose that the bits stored are associated with none of the values. Note that none of  $\{v_1, v_2, v_3\}$  are returnable from both servers. Now, imagine a single step, where the first server receives a message that contains  $v_2$  and, after the receipt of the message, stores  $v_1 + v_3$  at point  $P'$ . Then a reader which can access the bits stored in both servers can recover  $v_2$  by simply subtracting the contents of the two servers. However, the number of bits stored in the servers did not change in the step! Thus the proof method of [23] would not be generally applicable to algorithm  $A$ .

The proof technique of [23], in fact, applies when versions are encoded *separately*. For instance, if a server that receives information from three values  $v_1, v_2, v_3$ , and stores information of the form  $(f_1(v_1), f_2(v_2), f_3(v_3))$ , where  $f_1, f_2, f_3$  are three arbitrary functions. In this instance, every stored bit can be uniquely associated with a write operation.

It must be noted that the algorithm  $A$  considered in this section is hypothetical. We do not know whether we can construct meaningful algorithms that can encode across different versions. Our storage cost lower bound of Theorem 6.5 suggests that even if such an algorithm can be constructed, there would be little benefit from a storage cost perspective if the writers send values in a single phase.

## B A Simple Information-theoretic Lower Bound

In this section, we provide a simple information-theoretic lower bound on the storage cost incurred by any distributed shared memory emulation algorithm. We describe the lower bound in Theorem B.1. The theorem leads to lower bounds on the max- and total-storage costs, which are stated in Corollary B.2. After stating Theorem B.1 and Corollary B.2, we provide an informal description of the proof of Theorem B.1, followed by a formal description.

### B.1 Statement of Theorem B.1

**Theorem B.1.** *Let  $A$  be a single-writer single-reader shared memory emulation algorithm that implements a regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose that, in algorithm  $A$ , every server's state belongs to a set  $\mathcal{S}$ . Suppose that the algorithm  $A$  satisfies the following liveness property:*

*In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$ ,  $f \geq 1$ , then every operation invoked at a non-failing client terminates.*

*Then, for every subset  $\mathcal{N} \subset \{1, 2, \dots, N\}$  where  $|\mathcal{N}| = N - f$ ,*

$$\sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_n| \geq \log_2 |\mathcal{V}|.$$

The above theorem naturally implies a bound on the total- and max-storage costs as demonstrated in the following corollary.

**Corollary B.2.** *Let  $A$  be a single-writer-single-reader shared memory emulation algorithm that implements a regular read-write object whose values come from a finite set  $\mathcal{V}$ . Suppose that every server's state belongs to a set  $\mathcal{S}$  in algorithm  $A$ . Suppose that the algorithm  $A$  satisfies the following liveness property:*

*In a fair execution of  $A$ , if the number of server failures is no bigger than  $f$ ,  $f \geq 1$ , then every operation invoked at a non-failing client terminates.*

*Then*

$$\begin{aligned} \text{MaxStorage}(A) &\geq \frac{\log_2 |\mathcal{V}|}{N - f}, \text{ and} \\ \text{TotalStorage}(A) &\geq \frac{N \log_2 |\mathcal{V}|}{N - f}. \end{aligned}$$

*Proof of Corollary B.2.* We assume, without loss of generality, that  $|\mathcal{S}_1| \leq |\mathcal{S}_2| \leq \dots \leq |\mathcal{S}_N|$ . From Theorem B.1, we have

$$\sum_{n=1}^{N-f} \log_2 |\mathcal{S}_n| \geq \log_2 |\mathcal{V}|.$$

As a consequence, we have  $\log_2 |\mathcal{S}_{N-f}| \geq \frac{\log_2 |\mathcal{V}|}{N-f}$ . Therefore, we have  $\max_{n \in \{1, 2, \dots, N\}} \log_2 |\mathcal{S}_n| \geq \log_2 |\mathcal{S}_{N-f}| \geq \frac{\log_2 |\mathcal{V}|}{N-f}$ . Furthermore, we have  $\log_2 |\mathcal{S}_n| \geq \frac{\log_2 |\mathcal{V}|}{N-f}$  for every  $n \in \{N-f+1, \dots, N\}$ . This implies the following chain of relations.

$$\begin{aligned} \sum_{n=1}^N \log_2 |\mathcal{S}_n| &\geq \log_2 |\mathcal{V}| + \sum_{n=N-f+1}^N \log_2 |\mathcal{S}_n| \\ &\geq \log_2 |\mathcal{V}| + f \frac{\log_2 |\mathcal{V}|}{N-f} = \frac{N \log_2 |\mathcal{V}|}{N-f} \end{aligned}$$

This completes the proof.  $\square$

We now prove Theorem B.1.

## B.2 Informal Proof Sketch for Theorem B.1

Intuitively, the above theorem can be understood as follows. Consider any subset  $\mathcal{N} \subseteq \{1, 2, \dots, N\}$  where  $|\mathcal{N}| = N - f$ . Consider an execution of the algorithm  $A$  where servers  $\{1, 2, \dots, N\} - \mathcal{N}$  fail at the beginning of the execution. After the servers fail, a writer writes value  $v$  in the system and terminates. Because the algorithm is regular, any reader that begins after the termination of the write must recover the last written value  $v$  from the  $N - f$  servers in  $\mathcal{N}$ . Since the state of server  $i$  belongs to  $\mathcal{S}_i$ , the total number of possible configurations of the states of the  $N - f$  servers in  $\mathcal{N}$  is  $\prod_{i=1}^{N-f} |\mathcal{S}_i|$ . Since the value  $v$  can be any element of the set  $\mathcal{V}$  and the reader must recover the value through messages exchanged with the servers, there must be a one-to-one mapping from the set of values to the set of server states. Therefore, we need  $\prod_{n \in \mathcal{N}} |\mathcal{S}_n| \geq |\mathcal{V}|$ , which implies the result of Theorem B.1. We provide a formal proof below.

## B.3 Formal Proof of Theorem B.1

*Proof of Theorem B.1.* Consider any subset  $\mathcal{N} \subseteq \{1, 2, \dots, N\}$  where  $|\mathcal{N}| = N - f$ . We construct  $|\mathcal{V}|$  executions of the algorithm. In particular, for every value  $v$  in  $\mathcal{V}$ , we construct an execution  $\alpha^{(v)}$  of the algorithm as follows. In  $\alpha^{(v)}$ , the  $f$  servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  fail at the beginning of the execution. The servers in  $\mathcal{N}$  do not fail in  $\alpha^{(v)}$ . After the  $f$  servers fail, a write operation with value  $v$  begins and all components take turns in a fair manner until the write operation terminates. Since, in a fair execution of algorithm  $A$  where the number of server failures is at most  $f$ , any operation invoked at a non-failing client eventually terminates, we can ensure that the execution can be extended until the write terminates. Let  $\tilde{P}^{(v)}$  be some point after the termination of the write. At  $\tilde{P}^{(v)}$ , all the channels in the system act, delivering all their messages. Let  $P^{(v)}$  be some point in  $\alpha^{(v)}$  after the channels deliver their messages. At  $P^{(v)}$ , the write client fails. At some point after  $P^{(v)}$ , a read operation begins and all the components in the system take turns in a fair manner until the read terminates. Because the read client does not fail, and because the number of server failures is  $f$ , the read operation terminates in  $\alpha^{(v)}$ . The execution  $\alpha^{(v)}$  ends after the completion of the read operation.

For  $j \in \mathcal{N}$  denote the state of server  $j$  at point  $P^{(v)}$ . We denote by  $\vec{S}^{(v)}$ , the tuple  $(S_{j_1}^{(v)}, S_{j_2}^{(v)}, \dots, S_{j_{N-f}}^{(v)})$  of server states, where  $\mathcal{N} = \{j_1, j_2, \dots, j_{N-f}\}$ , and  $j_1 < j_2 < \dots < j_{N-f}$ . Note that  $\vec{S}^{(v)}$  is an element from the set  $\prod_{n \in \mathcal{N}} \mathcal{S}_n$ . To complete the proof of the lemma, it suffices to show that for two distinct values  $v, v'$  in  $\mathcal{V}$ , we have  $\vec{S}^{(v)} \neq \vec{S}^{(v')}$ . This is because  $\vec{S}^{(v)} \neq \vec{S}^{(v')}$  implies that there are at least  $|\mathcal{V}|$  elements in  $\prod_{n \in \mathcal{N}} \mathcal{S}_n$ , which implies the theorem statement owing to the following chain of relations:

$$\begin{aligned} \prod_{n \in \mathcal{N}} |\mathcal{S}_n| &\geq |\mathcal{V}| \\ \Rightarrow \sum_{n \in \mathcal{N}} \log_2 |\mathcal{S}_n| &\geq \log_2 |\mathcal{V}| \end{aligned}$$

So, to complete the proof, it is enough to show that for distinct values  $v \neq v'$ , we have  $\vec{S}^{(v)} \neq \vec{S}^{(v')}$ .

Assume for contradiction that there exist two different values  $v \neq v'$  such that  $\vec{S}^{(v)} = \vec{S}^{(v')}$ . We create an execution  $\beta$  of the algorithm where the operations are not regular, which would contradict the assumption that the algorithm is regular and complete the proof. The steps of  $\beta$  are identical to the steps of execution  $\alpha^{(v)}$  until the point  $P^{(v)}$  in  $\alpha^{(v)}$ . Consider the composite automaton that includes the servers, the readers, the channels amongst the servers, and the channels between the readers and the servers. The state of every component of this composite automaton at point  $P^{(v)}$  in  $\alpha^{(v)}$  is the same as the state of the corresponding component at point  $P^{(v')}$  in  $\alpha^{(v')}$ . This is because at both  $P^{(v)}$  and  $P^{(v')}$ , all the channels are empty, the servers in  $\{1, 2, \dots, N\} - \mathcal{N}$  have failed, and the state of any server from  $\mathcal{N}$  is the corresponding element of  $\vec{S}^{(v)}$ , which is equal to  $\vec{S}^{(v')}$ . In  $\beta$ , after the point  $P^{(v)}$ , all the components follow the steps of the execution

$\alpha^{(v')}$ . Clearly  $\beta$  is an execution of the algorithm  $A$ . Because  $P^{(v)}$  is after the termination of the write operation that wrote value  $v$  in  $\alpha^{(v)}$ , and because  $\beta$  is identical to  $\alpha^{(v)}$  until the point  $P^{(v)}$ , we infer that a write operation wrote value  $v$  in  $\beta$ . Because a read operation begins after  $P^{(v')}$  and returns  $v'$  in  $\alpha^{(v')}$ , and because every component follows the same steps in execution  $\beta$  as in execution  $\alpha^{(v')}$  after point  $P^{(v')}$ , we infer that a read operation begins in  $\beta$  after the termination of the write and returns value  $v'$ . If  $\beta$  is regular, this read operation is serialized after the write operation in  $\beta$ . Therefore, the read operation should return  $v$ . However it returns  $v'$  which is not equal to  $v$ . Therefore  $\beta$  is not regular, which contradicts the assumption that the algorithm is regular. Therefore, for any two different values  $v \neq v'$ , we have  $\vec{S}^{(v)} \neq \vec{S}^{(v')}$ . This completes the proof.  $\square$